EDN®

VOICE OF THE ENGINEER

DEC 2

Issue 23/2010

www.edn.com

USE GPUs TO BOOST ACCELERATION

DESIGNING RELIABLE POWER-LINE COMMUNICATIONS

SOC-PLL DESIGN REQUIRES TRADE-OFFS

Integrated Device Technology – Leader in PCI Express® Solutions

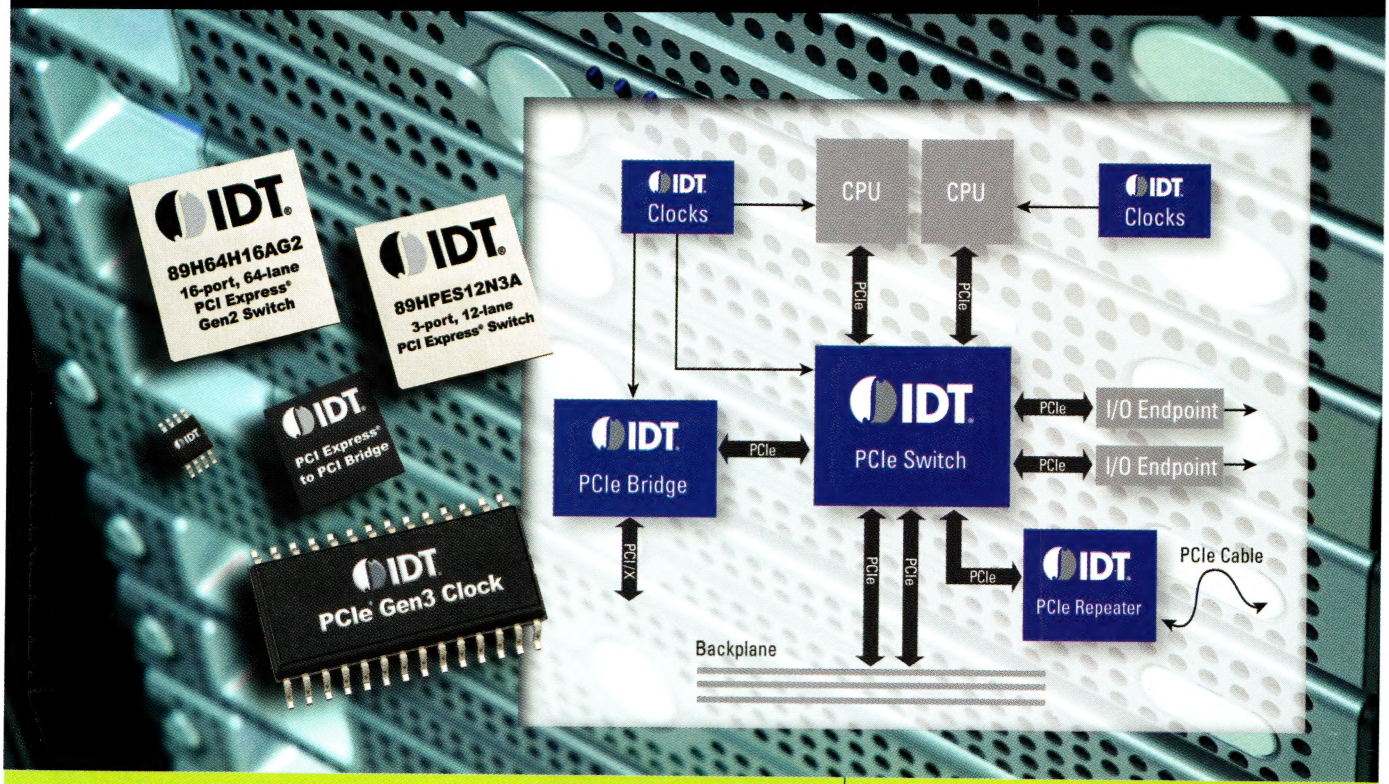

PCI Express Switches
Industry's most comprehensive family of high-performance, scalable PCIe switching solutions

- Gen 1 and Gen 2 Portfolio
- Up to 64 lane and 24 port devices
- Highly flexible port configurations
- Unprecedented 8 non-transparent bridging (NTB) functions to enable multi-root applications

PCI Express Repeaters
5Gbps Gen 2 PCIe Signal Repeaters

- Four and eight channel signal repeaters
- Compensates for cable and PCB trace attenuation and ISI jitter
- Programmable receiver equalization to 30dB and de-emphasis to -8.5dB
- Leading edge power minimization in active and shutdown modes
- No external bias resistors or reference clocks required
- Channel mux mode, demux mode, 1 to 2 channels multicast, and Z-switch function mode

PCI Express Bridges
High performance PCIe bridging to legacy PCI and PCI-X protocols

- Ultra low power version for consumer applications
- Forward mode buffer optimization
- Only PCI Express bridges with Short Term Caching for significant improvement in PCI Read performance
- Pin compatibility with competitive offering for dual sourcing solution

PCI Express Clocking Solutions
Industry's broadest offering of PCIe Gen1, Gen2 and Gen3 clock generation and buffering solutions

- World's First PCI Express Gen3 Family of Timing Devices
- Clock Generators and Synthesizers
- Spread Spectrum Clock Generators
- Fanout Buffers and Muxes
- Jitter Attenuators
- Frequency Translators
- Zero Delay Buffers

IDT
Integrated Device Technology

The Analog and Digital Company™

For more information visit us online at www.IDT.com

MachXO2™ PLD

MachXO2.
DO-IT-ALL DEVICES DEMAND A DO-IT-ALL PLD.

With the MachXO2 PLD family, you now have everything you need to meet your low power, small form factor, high functionality and low price requirements for your next consumer or system design.

Access free reference designs and software, development kits and supporting literature at:
latticesemi.com/machxo2

Applications

- Smart Phone
- Digital Camera
- GPS
- Server
- Router
- Base Station

And more...

Benefits

- Instant-on, non-volatile
- Flexible logic from 256 LUTs to 7,000 LUTs
- Embedded memory up to 240 Kbits
- Ultra low power as low as 19µW
- On-chip user flash memory & I^2C, SPI, timer/counter
- Prices starting as low as $0.75

And more...

Lattice® Semiconductor Corporation

latticesemi.com/machxo2

EDN 12.2.10 contents

Use GPUs to boost acceleration

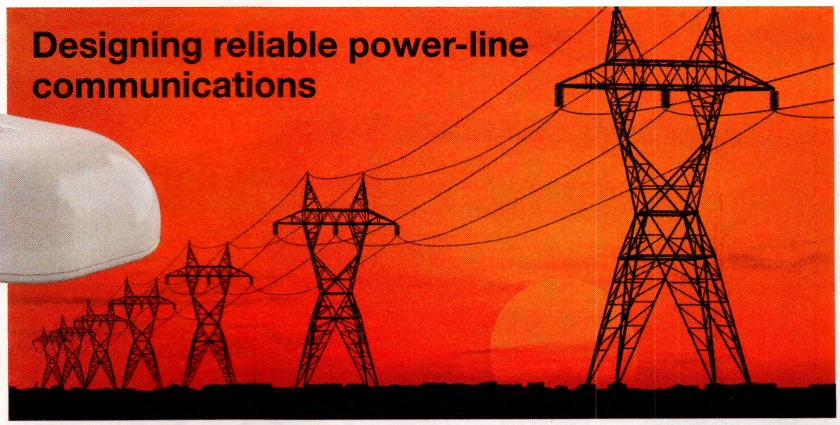

Designing reliable power-line communications

pulse

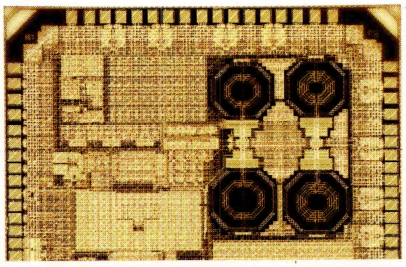

SOC-PLL design requires trade-offs

DESIGNIDEAS

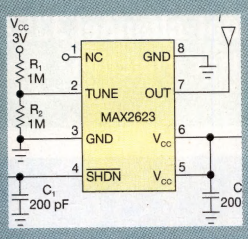

▶ Send your Design Ideas to edndesignideas@ubm.com.

Go ahead: define your most exacting waveform.

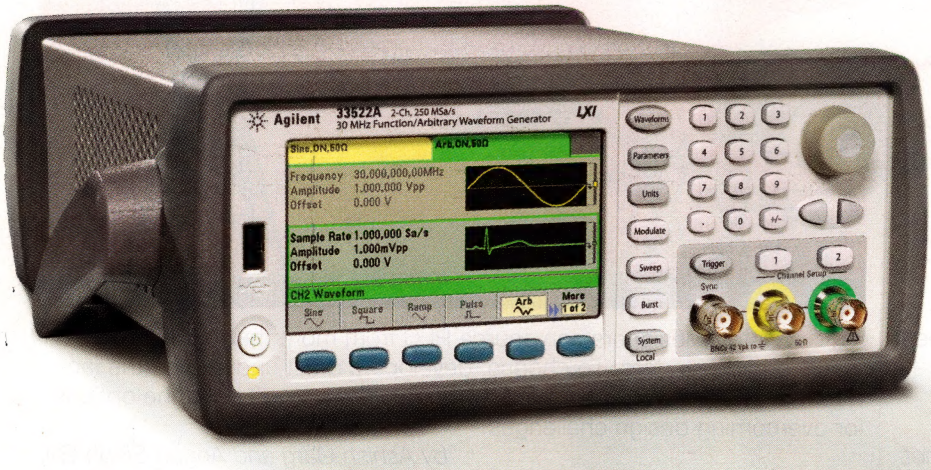

Generate a most exacting output.

You're developing new technologies every day that are faster, more efficient and imaginative. Agilent 33500A Series function/arbitrary waveform generators provide your waveforms with the highest signal fidelity. With 10x better jitter than anything in its class you have unparalleled control of signal frequency. And new breakthrough technology provides an unprecedented ability to generate arbitrary waveforms. That's confidence. That's Agilent.

33521A and 33522A
30 MHz sine, square, pulse bandwidth
250 MSa/s, 16-bit sampling
Point-by-point arbitrary waveforms
Dual-channel coupling and tracking

Agilent and our Distributor Network
Right Instrument.
Right Expertise.
Delivered Right Now.

TESTEQUITY
800-732-3457
testequity.com/agilent

Get 10% off the dual-channel 33522A
testequity.com/Agilent_Function

Agilent Technologies

contents 12.2.10

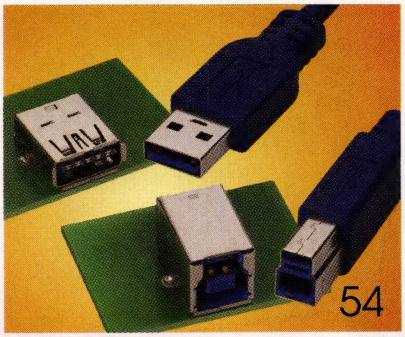

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out this Web-exclusive article:

NXP CEO: Capacity a challenge, but opportunity abounds
With opportunity, there is always challenge, and NXP Semiconductors is not immune to that fact.
→www.edn.com/101202toca

Channel relationships in focus

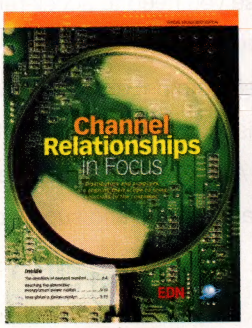

This *EDN*/NEDA special section takes a look at how suppliers and distributors continue to fine-tune their relationships in response to customer demands.
→www.edn.com/101202tocb

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. E-mail: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2.

Accelerating Your Success.™

BELL MICRO

AVNET Embedded

Embedded Solutions

Three Times the **Power**

People. Products. Services. The powerful combination of Avnet and Bell Microproducts provides the expertise you need to accelerate your success. Our combined team gives you access to world-class resources. Bringing industry leading line cards together, we now deliver the most extensive inventory of brand name systems, embedded hardware, displays, storage and software. And, with our enhanced services you have access to Avnet's leading ISO integration centers, financial solutions and supply chain strategies.

AVNET Embedded

50 Years AVNET ON THE NEW YORK STOCK EXCHANGE NYSE.

Supplier Authorized Distributor

1 800 332 8638

www.em.avnet.com/embedded

Follow us on Twitter!
www.twitter.com/avnetdesignwire

SUZANNE DEFFREE, MANAGING EDITOR, NEWS

Moderate but not disappointing growth years to follow record 2010 sales

You can almost hear the trepidation in chief executive officers' voices when the topic of 2011 growth comes up—not because 2011 won't show growth. The SIA (Semiconductor Industry Association) last month estimated 6% semiconductor-sales growth for next year, an estimate close to the 5.1% projection that iSuppli made in October.

Those figures represent good, solid, moderate growth, and they should cause no fear or anxiety to chief executive officers, chief financial officers, or vice presidents when they look to the year ahead. The quiet concern of these officials who report on their companies' sales to boards, Wall Street analysts, and stockholders arises because industry analysts will inevitably compare that 5 to 6% growth unfavorably with 2010's record growth.

In each 2010 quarter, many semiconductor-industry companies have shown exceptional sales. For example, Atmel, International Rectifier, and Microchip just last month each reported robust results that outdid analyst estimates. They are not alone, and many companies' third-quarter sales reports included phrases such as "record growth" and "beat expectations."

As we usher in 2011, though, sales growth may seem to fall as quickly as the Times Square ball. It's as undeniable as a New Year's Day hangover: Sales will be more moderate in 2011 than they were in 2010 for most semiconductor companies. The phrases noted here will appear less frequently in earnings reports. Inventory should move back in line with demand as the soft-landing recovery we are experiencing continues.

Some shortsighted industry participants on Wall Street or at analysis companies will reflect on that fact and blast negativity out to the financial community. They'll do so with the stench of 2009 still lingering in the air.

However, a big difference exists between moderate growth and disappointing growth or, for that matter, sales declines. No semiconductor-market-research company that I've come across so far has predicted a sales decline for 2011. Instead, they've predicted more moderate growth than the industry experienced in 2010.

Judging by research expectations, total sales will not be disappointing as we move into 2011. Slow and steady wins the race; I'd rather see stable 5 to 10% growth year over year from the industry *EDN* covers than the 32% we're expected to record this year. If you're gaining 32%, you're either coming off a horrific year or you've made some tremendous innovation or change, which would lead to another set of questions.

I wouldn't want to be in a chief executive officer's shoes next year. Their companies are working through the economic downturn and recovering, but shortsighted industry pundits may take a different view. Over the last two years, if it weren't for the shortsighted—those who bought homes without regard for increases in the cost of living, sold stocks without a second thought, and froze R&D and talent investment—the economy wouldn't be in the mess it's in to begin with.

It's as undeniable as a New Year's Day hangover: Sales will be more moderate in 2011 than they were in 2010 for most semiconductor companies.

Semiconductor-sales growth should continue through 2012, according to the SIA, which estimates an increase of 3.4% over 2011. This growth would put sales at about $330 billion in 2012. As such, the SIA's projected compound-annual-growth rate should be 13.4% for 2009 through 2012. That rate is not disappointing when you take a more reasonable, longer-term look at the semiconductor segment. EDN

Contact me at suzanne.deffree@ubm.com.

Can You Solve This?

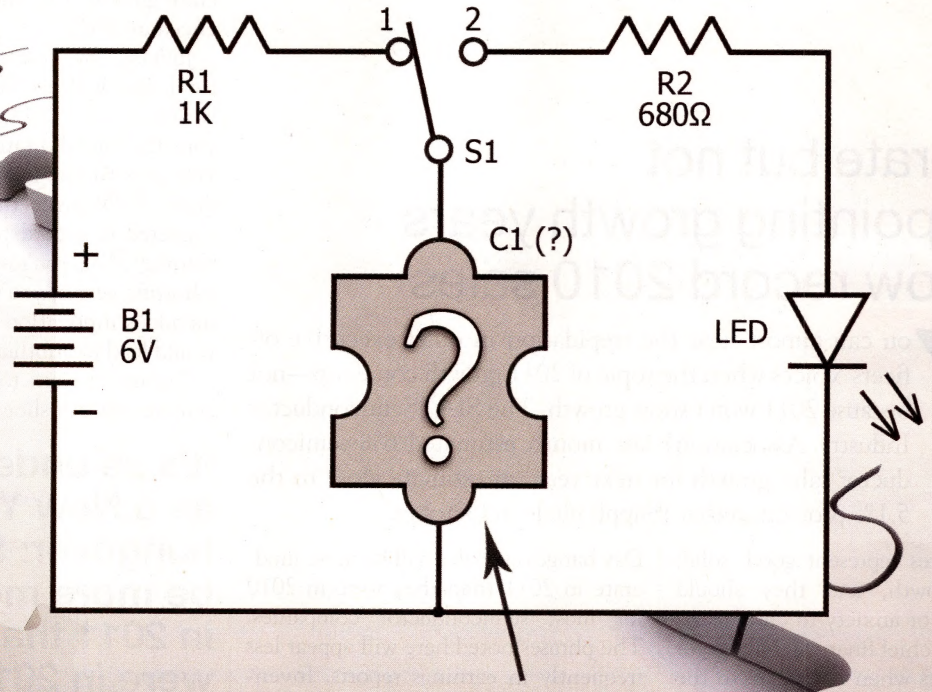

What is the missing component?

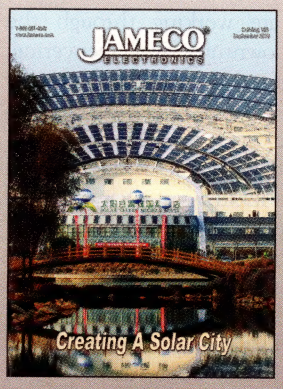

Electronics instructor Ollie Circuits planned to show his class of freshman electrical engineering students how to use a super capacitor as a memory back-up capacitor, but first he wanted to show how the students could make their own super capacitor and demonstrate its charge/discharge cycles with the simple circuit above. Most of the components were already on his workbench, the homemade super capacitor would be made from several layers of lemon juice-soaked paper towels interleaved between several layers of a mystery material to form a multi-layer stack. The stacked layers would then be sandwiched between the two copper-clad PC boards and held together with a rubber band. Ollie rushed to a nearby pet shop. What did he buy? Go to www.Jameco.com/brainpower to see if you are correct and while you are there, sign-up for our free full-color catalog.

1-800-831-4242 | www.Jameco.com

JAMECO® ELECTRONICS

VALUE POWER SUPPLIES BY MEAN WELL

Jameco is a proud distributor of Mean Well's full line of high-quality power supplies.

- Value pricing
- Inventory in stock
- Full line distributor
- Local technical support

Let us help you design Mean Well into your solution.

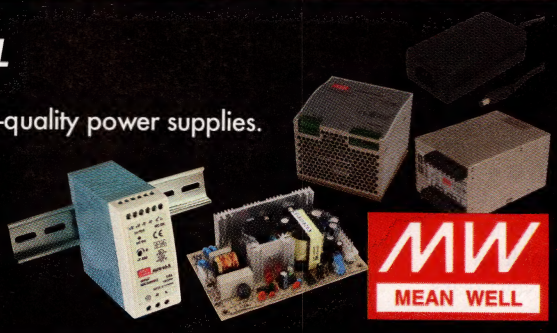

EDN

PUBLISHER, *EDN* WORLDWIDE
Russell E Pratt, 1-781-869-7982;
russell.pratt@ubm.com

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Judy Hayes, 1-925-736-7617;
judy.hayes@ubm.com

EDITOR-IN-CHIEF, *EDN* WORLDWIDE
Rick Nelson
Test and Measurement, DFx
1-781-869-7970;
richard.nelson@ubm.com

MANAGING EDITOR
Amy Norcross
Contributed technical articles
1-781-869-7971;
amy.norcross@ubm.com

MANAGING EDITOR—NEWS
Suzanne Deffree
Electronic Business, Distribution
1-631-266-3433;
suzanne.deffree@ubm.com

SENIOR TECHNICAL EDITOR
Brian Dipert
Consumer Electronics, Multimedia, PCs, Mass Storage
1-916-548-1225;
brian.dipert@ubm.com

TECHNICAL EDITOR
Margery Conner
Power Sources, Components, Green Engineering
1-805-461-8242;
margery.conner@ubm.com

TECHNICAL EDITOR
Mike Demler
EDA, IC Design and Application
1-408-384-8336;
michael.demler@ubm.com

TECHNICAL EDITOR
Paul Rako
Analog, RF, PCB Design
1-408-745-1994;
paul.rako@ubm.com

DESIGN IDEAS EDITOR
Martin Rowe,
Senior Technical Editor,
Test & Measurement World
edndesignideas@ubm.com

SENIOR ASSOCIATE EDITOR
Frances T Granville, 1-781-869-7969;
frances.granville@ubm.com

ASSOCIATE EDITOR
Jessica MacNeil, 1-781-869-7983;
jessica.macneil@ubm.com

CONSULTING EDITOR
Jim Williams,
Staff Scientist, Linear Technology
edn.editor@ubm.com

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg,
strassbergedn@att.net
Nicholas Cravotta,
editor@nicholascravotta.com
Robert Cravotta
robert.cravotta@embeddedinsights.com

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap
Kevin C Craig, PhD, Marquette University

LEAD ART DIRECTOR
Marco Aguilera

ASSOCIATE ART DIRECTOR
Tim Burns

PRODUCTION
Michael Ciardiello,
Director of Premedia Technologies
Jeff Tade,
Production Director
Brian Wu,
Publications Production Manager
Jeff Polman, Derric Treece,
Senior Production Artists
William Baughman, Lucia Corona
Ricardo Esparza,
Production Artists

EDN EUROPE
Graham Prophet, Editor, Reed Publishing
gprophet@reedbusiness.fr

EDN ASIA
Wai-Chun Chen, Group Publisher, Asia
waichun.chen@ubm.com
Kirtimaya Varma, Editor-in-Chief
kirti.varma@ubm.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
william.zhang@ubm.com
Jeff Lu, Executive Editor
jeff.lu@ubm.com

EDN JAPAN
Katsuya Watanabe, Publisher
katsuya.watanabe@ubm.com
Ken Amemoto, Editor-in-Chief
ken.amemoto@ubm.com

EXECUTIVE OFFICERS
Paul Miller,
Chief Executive Officer
Fred Gysi,
Chief Financial Officer
Mike Deering,
Chief Technology Officer
Ron Wall,
Senior Vice President, Publications
Kevin O'Keefe,
Senior Vice President, Events Division
Roger Burg,
Vice President, Operations
Jason Brown,
Vice President, E-media

UBM Electronics

EDN, 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com**. **Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change.

UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

pulse

INNOVATIONS & INNOVATORS

Motion sensor packs in three-axis gyro and accelerometer along with hardware accelerator

As 3-D-motion-sensing technology trends toward must-have status for consumer devices, MEMS (microelectromechanical-system)-based accelerometers, gyroscopes, and compasses have become more dominant components, all jostling for board space and processor and internal I^2C (inter-integrated-circuit)-bus bandwidth. Jumping on that bandwagon, Invensense has introduced the MPU-6000 motion processor, which features a three-axis gyroscope, a three-axis accelerometer, and a Digital Motion Processor hardware accelerator for processing nine-axis motion algorithms, all on the same silicon die, in a 4×4×0.9-mm package.

Six- and nine-axis motion sensing brings improved and new capabilities, such as precise sensing of hand jitter to improve image quality and video stability; GPS (global-positioning-system) dead reckoning for vehicles and indoor pedestrian-navigation and motion-based user interfaces; and augmented reality for immersive-gaming experiences. The MPU-6000 also allows users to input data from a three-axis compass that then enables nine-axis sensor processing.

The MPU-6000 includes a range of dynamic, full-scale capabilities at ±250, ±500, ±1000, and ±2000°/sec for angular-rate sensing and ±2, ±4, ±8, and ±16g for linear acceleration sensing, allowing the device to work in motion applications that include slow-motion menu selection as well as fast hand gestures, all with 16-bit resolution. Rate-noise performance for image stabilization, pointing, and gaming applications is 0.005°/sec/$\sqrt{Hz}$. High-accuracy factory calibration of ±1% initial sensitivity reduces customers' calibration requirements.

The gyroscope operates at a resonant frequency higher than 27 kHz, making the MPU-6000 immune to interference from audible frequencies of 20 Hz to 20 kHz, such as music, phone ringers, crowds, and white noise. This immunity is critical in noise-sensitive applications, such as image stabilization.

Other features include on-chip16-bit ADCs, programmable digital filters, a precision clock with 2% accuracy over –40 to +85°C, an embedded temperature sensor, programmable interrupts, and 5.5-mA current consumption. The device fits into a 24-pin QFN package and is available with an I^2C interface and an SPI (serial-peripheral interface), an operating-voltage range of 2.5 to 3.6V, and a logic-interface voltage of 1.71 to 3.6V and sells for less than $3. **—by Margery Conner**

Invensense, www.invensense.com.

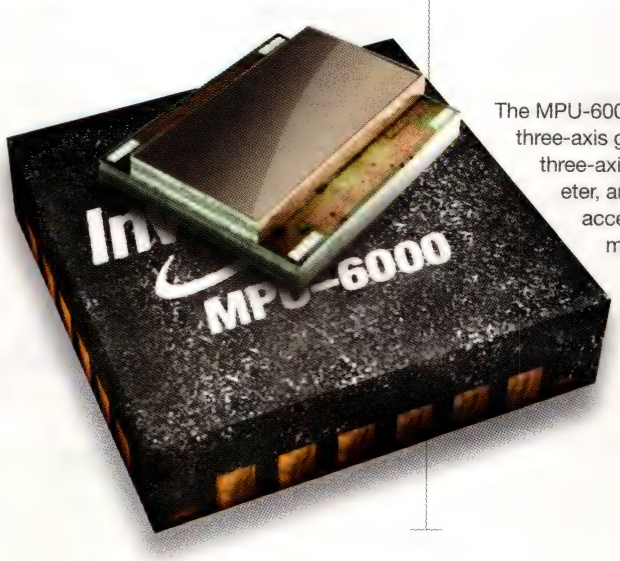

The MPU-6000 combines a three-axis gyroscope, a three-axis accelerometer, and a hardware accelerator for motion-processing algorithms, all on one die.

TALKBACK

"Am I missing something? (I could have a pizza, too, and not mess with any of this stuff. Of course, pizza costs less than $1.)"

–Technician Phil Stewart, in *EDN*'s Talkback section, at http://bit.ly/9ZcQoj. Add your comments.

Intel Embedded Solutions: In-Vehicle Infotainment

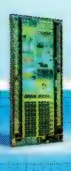

Performance headroom at under 5 watts.

Online goes on the road.

Intelligence where you need it. The Intel® Atom™ processor brings expanded capabilities to low-power applications—fueling an entirely new in-vehicle infotainment experience. Get rolling with the developer's playbook at **intel.com/embedded/intelligence**

intel

Sponsors of Tomorrow.

Ansys 13.0 gets fidelity and performance boost

Ansys Inc has launched the Ansys 13.0 engineering-simulation technology suite. The new release includes hundreds of features that provide performance boosts with respect to fidelity, productivity, and performance.

The product features an electromagnetic transient solver that produces higher-fidelity results in dynamic simulation environments. An adaptive architecture brings higher productivity through features that minimize the time and effort product-development teams invest in simulation, allowing users with different engineering specialties to work collaboratively to exchange data and develop real-world simulations that incorporate multiphysics.

The suite also achieves performance innovation through software and computational power. It provides faster speedup ratios than those of previous software releases, enabling complex multiphysics simulations to run more quickly and efficiently.

"Ansys 13.0 builds on the foundation of previous Ansys releases," says Jim Cashman, president and chief executive officer of Ansys. "It especially addresses user feedback by incorporating valuable capabilities that compress design cycles, optimize product performance across multiple physics, maximize the accuracy of virtual prototypes, and automate the simulation process."

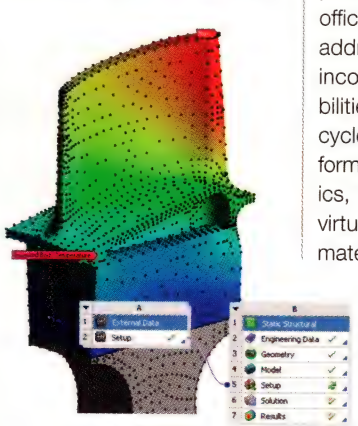

Ansys 13.0 has an external data mapper imports external data in the form of a column text file defining a point cloud and projects the data onto the current mesh. The feature allows users from different groups to straightforwardly exchange data.

Cashman expects the release to create a competitive advantage for users, making it easier and faster to bring innovative new products to market.

A customer ready to employ Ansys 13.0 is Red Bull Technology. "Red Bull Racing knows a lot about the need for speed," says Steve Nevey, business-development manager at Red Bull. "Ansys helps us move to the next level of racing by empowering us to innovate car designs much quicker, which makes us more competitive race by race."

> Engineers in structural mechanics can use 3-D rezoning for applications that involve large shape deformations.

Ansys reports that the new software release includes solver methods that leverage advanced technology to more quickly deliver accurate answers. For example, engineers in the structural-mechanics arena can use 3-D rezoning for applications that involve large shape deformations. It allows engineers to stop their simulation as their mesh becomes distorted and then remesh the material's current state and continue the simulation.

The Ansys HFSS (high-frequency-simulator-system) transient solver for dynamic electromagnetic simulations, such as broadband and radar, incorporates automatic adaptive mesh refinement and an innovative local time-stepping procedure to accurately represent the geometry and fields and to optimize runtime, stability, and efficiency. In addition, a new hybrid solver bridges finite-element and integral equation methods for high-frequency electromagnetic problems, making HFSS the first commercial code to have this capability, according to the company.

The new release employs the Ansys foundation of open and flexible architecture. Integration with Microsoft Excel enables users to interact with spreadsheets that contain analytic representations of models and parameter-table definitions.

Other enhancements include a cut-cell meshing feature that automatically produces nearly all hexahedral elements on complex 3-D geometry; an external data mapper that imports data in the form of a column text file defining a point cloud; and tighter integration among the Ansys Maxwell, HFSS, and mechanical solvers when doing electromagnetic–thermal–structural simulations.

Ansys 13.0's power enhancements include a new and, the company says, unique application of variational technology that reduces solution time by a factor of five to 10 for harmonic analysis. The new release also features greater support for GPUs (graphics-processing units), which can offload complex, time-consuming algorithms to increase processing speed and accuracy.

The integration of features new to release 13.0 aligns with the Ansys vision to make multiphysics engineering-simulation tools accessible to all who need them. Ansys 13.0 will be available for customer downloading this year. **—by Rick Nelson**

Ansys, www.ansys.com.

DILBERT By Scott Adams

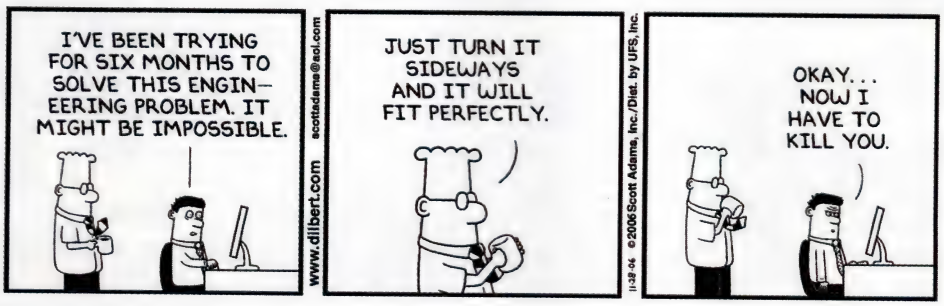

DENSE. FAST. FLEXIBLE.

SEARAY
OPEN PIN FIELD INTERCONNECTS

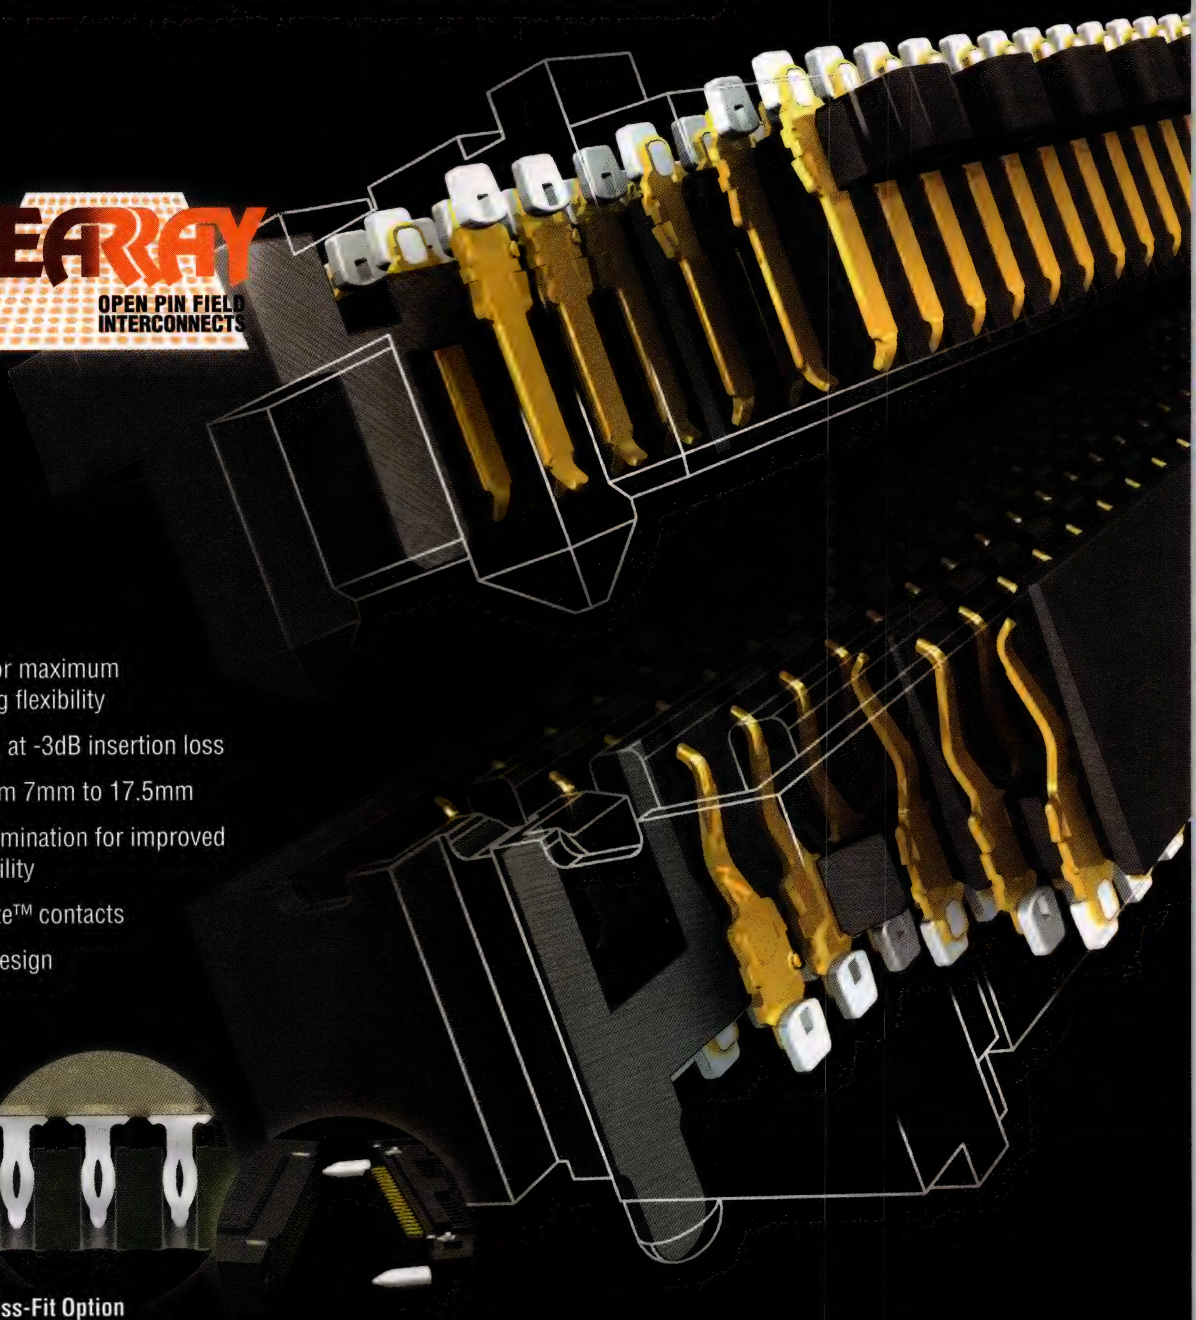

- Up to 500 I/O's for maximum grounding/routing flexibility
- Rated at 9.5 GHz, at -3dB insertion loss
- Stack heights from 7mm to 17.5mm
- Solder charge termination for improved processing reliability
- Rugged Edge Rate™ contacts
- 0,8mm pitch in design

Press-Fit Option

Right Angle Version

samtec

Transmission Line Solutions

www.samtec.com/searay

Equalizer IC receives video over one mile of Category 5 cable

Intersil's new ISL59605 MegaQ video-equalizer IC requires no active circuitry at the camera end, so you need not retrofit cameras mounted on buildings or poles. The ISL59605 compensates for high-frequency cable loss as high as 60 dB at 5 MHz and source-amplitude variations as high as ±3 dB. The device is compatible with both color and monochrome specifications of the NTSC (National Television System Committee) and PAL (phase-alternating-line) standards. It operates from a 5V supply; consumes only 500 mW; and has ac-coupled, internally dc-restored inputs. The output drives 2V p-p into two source-terminated 75Ω loads. The unit continuously monitors the cable characteristics and compensates for long-term cable aging. Several optional locking modes help maintain lock in noisy environments; an optional serial interface provides signal monitoring.

The ISL59605 is available in a 4×4-mm, 20-lead QFN package with suggested retail prices of $18.98 (1000). It operates over a –40 to +105°C temperature range. The ISL5960X-EvalZ demonstration board costs $250. The $350 ISL5960X-SPI (serial-peripheral-interface) EvalZ provides user access to SPI for signal monitoring.

—by Paul Rako

Intersil, www.intersil.com.

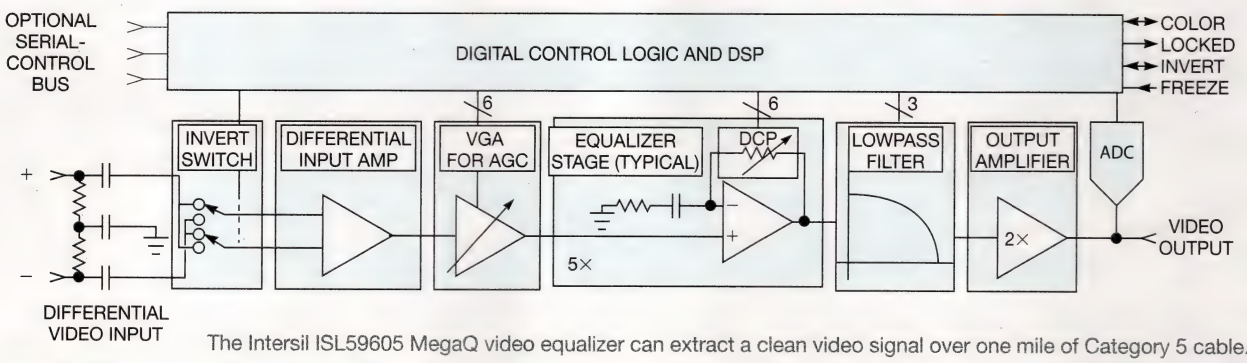

The Intersil ISL59605 MegaQ video equalizer can extract a clean video signal over one mile of Category 5 cable.

Solid-state-drive controllers get serious

SandForce recently released details on the initial four family members in its SF-2000 series of solid-state-drive controllers. The base SF-2000 controller beefs up its predecessors' system interface to the SATA (serial-advanced-technology-attachment) 6-Gbps variant, along with as many as 32 NCQ (native-command-queuing) entries. It also supports the newer Toggle and ONFI (Open NAND Flash Interface) 2 NAND options, which speed information transfers between the controller and the storage array. It can also activate—that is, simultaneously read from, write to, or both—two-times more flash-memory die than could the SF-1000- and SF-1500-controller generations.

This enhanced controller potential translates to as many as 60,000 IOPS (input/output operations/sec) for random read and write transfers of 4-kbyte data packets, along with speeds as high as 500 Mbytes/sec for sequential reads and writes. These estimates are for 34-nm MLC (multilevel-cell) flash memory, not for the inherently faster, albeit costlier on a per-bit basis, SLC (single-level-cell) devices.

The SF-2300 sells for approximately $50 (one, with quantity discounts available) and builds on the SF-2000 foundation with operation over the industrial-temperature range. The enterprise-targeted SF-2500 has a real-time AES (Advanced Encryption Standard)-256 encryption engine and double-encryption support—that is, the ability to configure unique passwords for different regions of the solid-state drive's total address range. And the SF-2600, which sells for $150 (single units), takes initial tentative steps toward full SAS (serial-attached SCSI) support.

The SF-2600 requires an external SAS-to-SATA bridge chip to comprehend the all-important SCSI (small-computer-system-interface) command set. Through hardware write-organization-alignment "hooks," however, it embeds support for non-512-byte sector variants, including 520-, 524-, and 528-byte and 4-kbyte sectors; the 4-kbyte sectors add a data-integrity field.

All SF-2000 family members use a common silicon design, with firmware variation the predominant means of differentiating products. And they all support EDAC (error detection and correction) at BCH (Bose/Chaudhuri/Hocquenghem) code levels as high as 55 bits versus the previous generations' 24-bit BCH-code capabilities, which SandForce brands as RAISE (redundant array of independent silicon elements), and commensurate with the controllers' cognizance of low-cost but low-reliability 3-bit-per-cell MLC-flash memory.

—by Brian Dipert

SandForce, www.sandforce.com.

12.02.10

VOICES

Kaufman Award winner Pat Pistilli: DAC and the birth of the EDA industry

Pat Pistilli is the 2010 recipient of the Phil Kaufman Award, which the DAC (Design Automation Conference) executive committee presents annually. He received the award for his pioneering efforts that led to the formation of the EDA industry and for creating and managing DAC since 1964. *EDN* spoke with Pistilli to get his perspective on the evolution of DAC and EDA. Read the full interview at www.edn.com/101202pa.

DAC has become the EDA industry's premier event. When you started, wasn't DAC purely a technical conference?

A For the first six years, the conference was known as SHARE [Society to Help Avoid Redundant Effort], and the technical papers were selected by me and a small group of industry CAD [computer-aided-design] people. It was invitation only. The first conference was in 1964, which we started with $1000 of our own money. We eventually got that money back.

What was the key for drawing attendees to DAC?

A When we started, I worked at Bell Labs, where we developed our own tools. The EDA industry developed later when people broke off from these internal groups to start their own companies. As the industry grew, companies were popping up, so there was motivation to add vendors to DAC. It became important to bring in exhibitors for smaller companies that could not afford to develop their own tools in-house.

We were an all-volunteer organization for 21 years. The executive-committee members were spending a lot of time on it, and it was time to bring in an outside vendor to run it. That is when my wife, Marie, persuaded me to leave Bell Labs to start MP Associates, and we continue as a vendor of DAC. Many people think we own DAC, but we don't, though Joe Costello [former chief executive officer of Cadence] once said, "The only one making money in EDA is Pistilli!"

Tell me about the Pat Pistilli Scholarship.

A During the volunteer years, we generated a surplus of $200,000 to $250,000. The first year that MP ran DAC, we generated a surplus of nearly $1 million. The surplus is smaller now, but it is used to sponsor the university booth and to provide travel grants to students. The scholarship was established for the advancement of computer science and electrical engineering, annually awarding two $20,000 scholarships to members of underrepresented groups, including women, minorities, and the disabled.

In some of your interviews, you have said that few new EDA products have come out over the past few years. Can you expand on that thought?

A There was a time that CAD tools exceeded manufacturing capability; now it is the reverse. Tools have not kept up with what can be built. Now, all the new stuff is being done at universities, funded by the EDA companies. We had 150 people in CAD development at Bell Labs. Now, the internal CAD departments are spending most of their time just coordinating tools they are purchasing, not developing.

Do you think the EDA industry does enough to make tools work together?

A The EDA industry is competitive. They are not going to help each other. Andy Graham [former president of the CFI (CAD Framework Initiative) and the Si2 (Silicon Integration Initiative)] tried and failed. Some standards are not accepted because the vendors have too much invested in proprietary formats. I wish it would happen; it would make it easier for people to buy tools and have them work together. Companies are looking out for their own interest, for the bottom line.

Who was the most interesting person that you have met in EDA?

A There is no question at all! It was Richard Newton. He was a good, good friend. He was a strict vegetarian, and Marie would make him a special dinner. We would finish off a bottle of wine or two and talk for hours. He was brilliant. A lot of companies had Richard on their boards. He was most respected.

You've been interviewed many times in recognition of receiving the Kaufman award. Is there one question that you think should have been asked that hasn't?

A The one thing nobody has asked me: "What would you do differently if you knew in 1964 what you know now?" I had a chance in the early 1970s to start my own CAD-development company. A DAC chairman asked me to leave Bell Labs. I couldn't do it because I was having so much fun developing these things. I would have been a multimillionaire now. If I had to do it over again, I may have jumped. But I'm not sorry I didn't. Bell Labs was so good to me for 21 years. Without it, DAC would not be what it is today. **—interview conducted and edited by Mike Demler**

LINEAR TECHNOLOGY

DESIGN NOTES

Complete Energy Utilization Improves Run Time of a Supercap Ride-Through Application by 40% – Design Note 485

George H. Barbehenn

Introduction

Many electronic systems require a local power source that allows them to ride through brief main power interruptions without shutting down. Some local power sources must be available to carry out a controlled shutdown if the main power input is abruptly removed.

A battery backup can supply power in the event of a mains shutdown, but batteries are not well suited to this particular application. Although batteries can store significant amounts of energy, they cannot deliver much power due to their significant source impedance. Also, batteries have finite lives of ~2 to 3 years, and the maintenance required for rechargeable batteries is substantial.

Supercapacitors are well suited to such ride-through applications. Their low source impedance allows them to supply significant power for a relatively short time, and they are considerably more reliable and durable than batteries.

Complete Energy Utilization Maximizes Run Time of Supercap Ride-Through Application

Figure 1 shows a complete 3.3V/200mA ride-through application that maximizes the amount of power extracted from the supercap to support the load.

The main components of the ride-through application include:

- The LTC®4425 complete 2A supercapacitor charger. It clamps the individual cell voltages to ensure that the cells do not overvoltage during charging and balances the cells throughout charge and discharge.
- The LTC3606 micropower buck regulator produces the regulated 3.3V output.
- The LTC4416 dual ideal diode switches the supercap in and out depending on need.

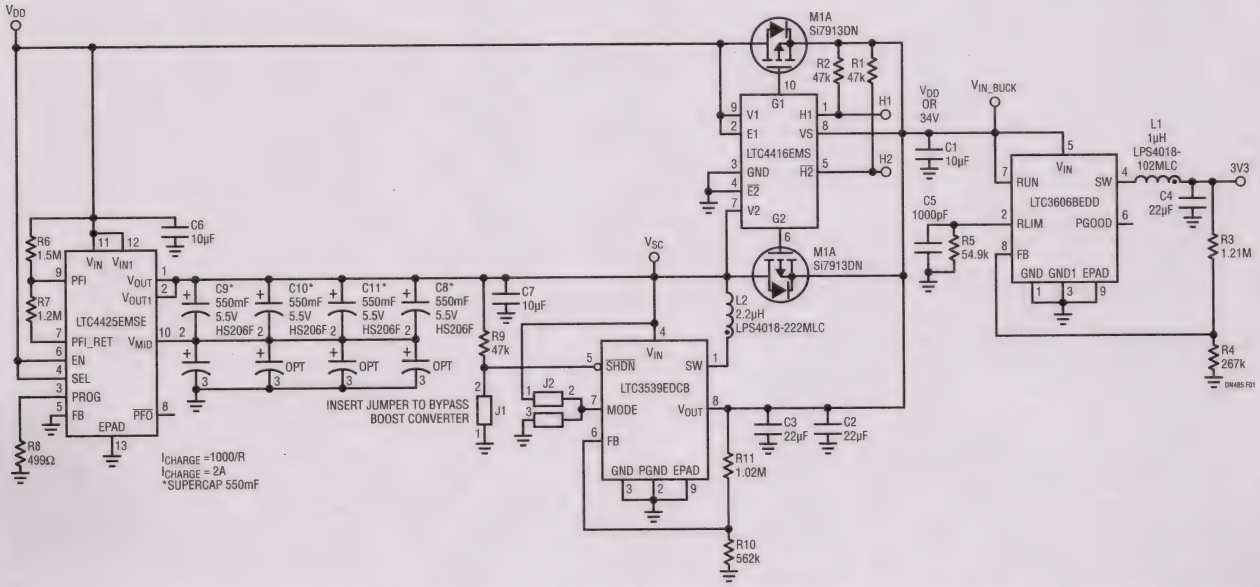

Figure 1. This Supercap-Based Power Ride-Through Circuit Maximizes Run Time Using an Energy Scavenging Scheme

- The LTC3539 micropower boost regulator with output disconnect recovers nearly all the energy in the supercap and it keeps the input to the LTC3606 above dropout as the supercap voltage drops. This boost regulator operates down to 0.5V.

40% Improvement in Run Time

Figure 2 shows the waveforms if the LTC3539 boost circuit is disabled. Run time from input power off to output regulator voltage dropping to 3V is 4.68 seconds. Figure 3 shows the waveforms if the LTC3539 boost circuit is operational. Run time from input power off to the output regulator dropping to 3V is 7.92 seconds. Note in Figure 3 that the output is a steady 3.3V voltage with a sharp cutoff.

How it Works

When the LTC3539 boost regulator is disabled, as soon as input power falls, the LTC4416 ideal diodes switch the input energy supply for the LTC3606 buck regulator to the supercap. In Figure 2, the voltage across the supercap (V_{SC}) is seen to linearly decrease due to the constant power load of 200mA at 3.3V on the buck regulator output (3V3).

In Figure 3, when the LTC3539 boost regulator is enabled, the voltage across the supercap (V_{SC}) is seen to linearly decrease due to the constant power load of 200mA at 3.3V on the buck regulator. When the voltage at V_{SC} reaches 3.4V, the regulation point of the boost regulator, the boost regulator begins switching. This shuts off the ideal diode and disconnects the buck regulator from the supercapacitor. The energy input to the buck regulator is now the boost regulator's output of 3.4V.

Because the input of the buck regulator remains at 3.4V, its output remains in regulation. When the boost regulator reaches its input UVLO and shuts off, its output immediately collapses, and the buck regulator shuts off.

Maximizing Usage of the Energy in the Supercap

Because each power conversion lowers the overall efficiency, the boost circuit should be held off as long as possible. Therefore, set the boost regulator output voltage as close to the buck regulator input dropout voltage as possible, in this case, 3.4V.

If the supercapacitor is initially charged to 5V, then the energy in the supercapacitor is 6.875J:

$$\frac{1}{2}CV^2 = \frac{1}{2}0.55F \bullet 5^2 = 6.875J$$

0.67W (3.33 • 0.2A)

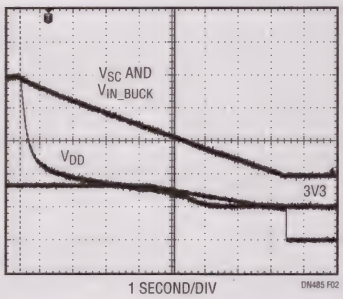

Figure 2. Power Ride-Through Application Results without Boost Circuit

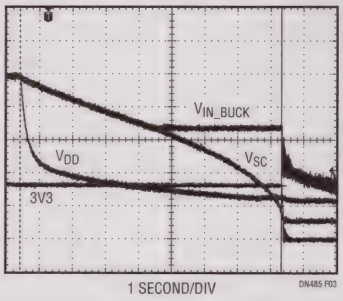

Figure 3. Power Ride-Through Application Results with Boost Circuit Enabled. The Boost Circuit Yields a 40% Improvement in Run Time

The output power is 3.33V • 0.2A = 0.67W, so the percentage of energy extracted from the full supercap when the boost regulator is disabled is 45.1%:

$$\frac{\varepsilon_{LOAD}}{\varepsilon_{CAP}} = \frac{0.67 \bullet 4.68s}{6.875} = 45.1\%$$

The percentage of the energy extracted from the supercap's available storage when the boost regulator is enabled is 77%:

$$\frac{\varepsilon_{LOAD}}{\varepsilon_{CAP}} = \frac{0.67 \bullet 7.92s}{6.875} = 77\%$$

This represents a 40% improvement in ride-through run time—significant when seconds count.

Conclusion

The run time of any given supercapacitor-based power ride-through system can be extended by 40% if energy is utilized from the discharging supercap. This is particularly relevant if the supercapacitor charge voltage is reduced to ensure high temperature reliability.

Data Sheet Download

www.linear.com

For applications help,
call (978) 656-3752

BY HOWARD JOHNSON, PhD

Impulsive behavior

As I advanced my cart through the check-out line at Walmart, the small child behind me wandered into the impulse zone filled with candy and pop magazines. At once, he commenced wailing and pleading to his mother. Apparently, this strategy had worked for him before. He continued relentlessly until she agreed to a small purchase. About that time, a pair of children in the next aisle moved to a similar position and set up a chorus of wails twice as grating as the first.

"As each child reached a certain point in line," I told my friend Chris "Breathe" Frue, "they became excited in the same way. That reaction forms the impulse response of this system. In reaction to a steady stream of children, the same reaction plays out over and over, in repetitive fashion.

"That must be excruciating for the check-out workers," said Breathe, a musician and audio technician who wants to learn more about equalizers.

"Consider what happens if the children arrive at irregular intervals or in groups," I said, drawing on a paper napkin the axes of a graph showing intensity versus time, with three curves at the top (**Figure 1**). "Suppose that the first child arrives, making a standard impulse response. That's the curve at the top left. Now, assume that the next two arrive together. If all children make the same standard response, the wails at that time will double, according to the number of participants. That phenomenon is superposition."

"Does the third curve then represent time invariance?" asked Breathe.

"It is certainly a good example of that," I replied. "Assume that each child reacts in turn according to when he or she reaches the candy display. If the last child is delayed, perhaps because she climbs out of the shopping cart and runs off, her ultimate reaction is then delayed by the exact time it takes for her mother to bring her back to the line, but that reaction remains otherwise unchanged. That phenomenon is time invariance." Drawing a final curve at the bottom, I said, "The composite response equals the sum of the three curves at the top."

"Are you saying that linear time-invariant systems react just as simply as children at Walmart?" he asked.

"I'm saying that if you stimulate any linear system with one short, intense pulse, the equivalent in this problem of one standard child, you see a response characteristic of that system. It's called the impulse response. You can then calculate the response to any future series of such short pulses as a sum of delayed and scaled copies of the one true impulse response."

"But real-life signals are not series of short pulses," Breathe said.

"But they are," I replied. "Imagine a train of short pulses, all packed shoulder to shoulder like a picket fence. In the limit, you can approximate any signal to within any degree of accuracy using a train of short enough pulses."

"Let's see if I have this straight," said Breathe. "Think of your input signal as a train of short pulses, each having an amplitude corresponding to the signal value at that point. Each pulse creates its own impulse response, scaled according to the pulse amplitude and delayed according to the pulse position in time. Superimpose all those little impulse responses, and you get the final system output."

"Astonishing, isn't it?" I said. "Any system's impulse-response waveform contains all the information about what a linear system can do. Measure that waveform, and you know all there is to know about the linear behavior of that system."EDN

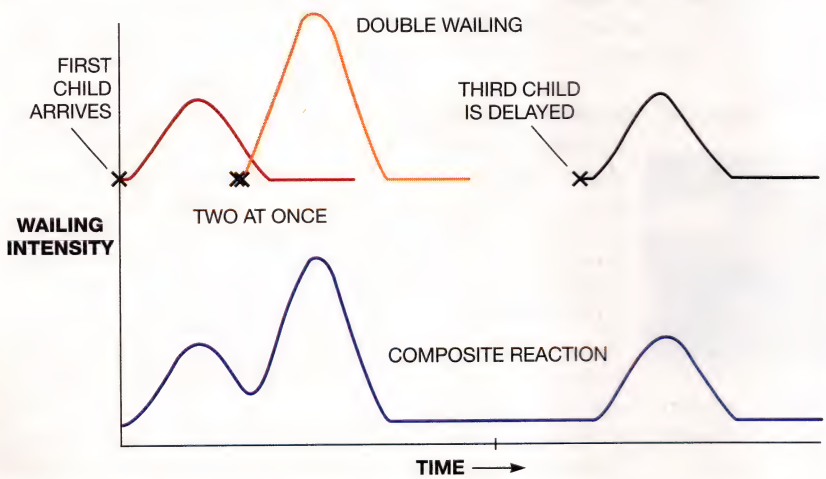

Figure 1 Each child generates a consistent impulse response.

Howard Johnson, PhD, frequently conducts technical workshops for digital engineers. Visit his Web site at www.sigcon.com.

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

Nanophotonics drive next-generation data rates

The field of nanophotonics is an area of growing importance in nanotechnology. Nanophotonics uses solid-state-created optical transmitters, waveguides, interconnects, and receivers for data signals between electronic circuits. These circuits may be other chips on the same substrate, on another module, or some distance away on another board or machine that connects through fiber-optic cable.

Most network data interfaces now use CMOS-driven copper interconnects. To drive data rates from the megahertz and gigahertz range to the hundreds of gigahertz and tetrahertz level, optical connections must replace the copper.

Long-reach applications, such as undersea cable and fiber to homes through central switching facilities, have long used optical interconnections. Standard-geometry semiconductor technology has made this approach possible for distances longer than 10m. Smaller and more efficient laser/detector pairs and high-data-rate electronics employing nanotechnology are necessary for addressing higher-speed technologies with shorter-range connections.

You can use standard CMOS processes to build these high-speed blocks with integrated photonics. The process involves standard CMOS transistors; a modulator; a waveguide, which acts as the optical interconnection on the die; and a diffraction grating. A key portion of the design is the laser-light source. The traditional method of achieving this source was to use lateral laser diodes.

High-density laser arrays and corresponding detectors now use VCSEL (vertical-cavity-surface-emitting-laser) devices. The laser in these devices comes out at a 90° angle to the surface of the device. The approach has the advantage of easy assembly in packaging in flip-chip packages and has narrow frequency dispersion, yielding high-efficiency single-mode performance (**Figure 1**).

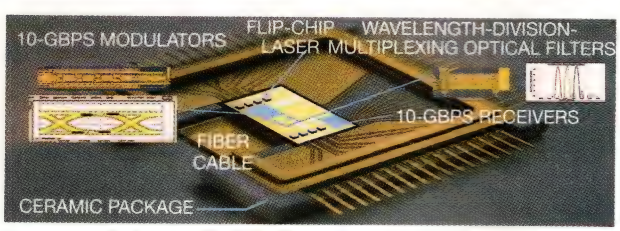

Figure 1 High-density laser arrays and corresponding detectors now use VCSEL devices, in which the laser exits at a 90° angle to the surface of the device. The approach allows easy assembly in flip-chip packages and has narrow frequency dispersion, yielding high-efficiency single-mode performance.

A high-speed electronic SERDES (serializer/deserializer) block is a key component of photonic interconnects and signal-processing circuitry. This block converts the parallel-bus-based data into serial streams so that an output block can convert them to high-speed photonic drivers. The output of these blocks drives solid-state laser transmitters; the receivers feed their input on the other end. Current SERDES devices target use in short-reach applications for on-chip and module-to-module connections and long-reach applications for traveling through more than 10m of fiber-optic cable. These blocks, fabricated in a 65-nm process, operate at 12 Gbps for short-reach applications and 20 Gbps for long-reach applications. Long-reach applications typically involve a large number of channels, so power is an issue. The 20-Gbps approach uses a 1.1V supply. Low-power applications using 0.95V supplies achieve 11 Gbps in blocks using a 32-nm process.

To drive data rates from the megahertz and gigahertz range to the hundreds of gigahertz and tetrahertz level, optical connections must replace the copper.

Current devices have as many as 576 differential-signal pairs and thus favor the VCSEL structure. Chips such as the Mellanox Infiniswitch IV support 36 40-Gbps ports, comprising four channels operating at 10 Gbps each, or 120-Gbps ports, comprising a group of 12 10-Gbps pairs. Other circuits employ a 100-Gbps channel comprising 10 10-Gbps pairs or eight 12.5-Gbps pairs.

Because of system power requirements, these applications were not possible until the emergence of nanoscale technology. A typical photonic switch would have 64 to 128 ports operating at 160 to 640 Gbps per port and consume less than 400 fJ/bit of I/O power. EDN

Pallab Chatterjee is vice chairman of the IEEE San Francisco Bay Area Nanotechnology Council. You can reach him at pallabc@siliconmap.net.

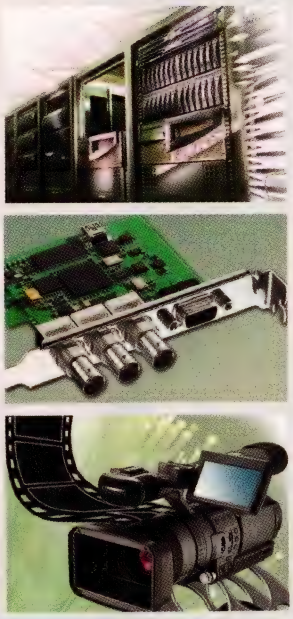

GET CLOCKS AND OSCILLATORS FAST

LOW-JITTER, HIGH-PERFORMANCE CUSTOM SAMPLES OVERNIGHT; PRODUCTION QUANTITIES IN 2–4 WEEKS

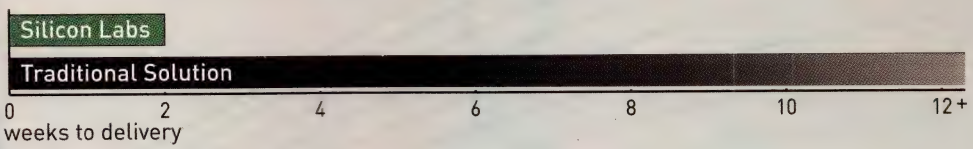

Silicon Labs' complete portfolio of industry-leading XOs, VCXOs, clock generators, jitter attenuating clocks and clock buffers set a new standard for flexibility, performance and lead time.

- Most frequency flexible timing solutions in the industry; ideal for multi-rate applications.
- Available in industry-standard, drop-in compatible packages.
- Based on patented DSPLL® and MultiSynth technologies, these low jitter products improve system performance, reduce BOM, minimize board space and simplify system design.
- All custom devices are available without NRE fees and <2-week lead times.

CUSTOM DEVICES WITHOUT NRE

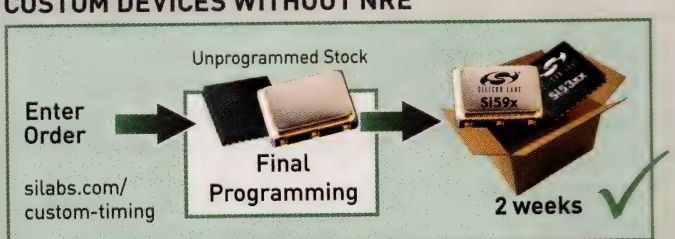

CLOCK GENERATORS

- Generates any frequency on any output up to 710 MHz
- Low phase jitter: <1 ps RMS
- Exact frequency synthesis (0 ppm error)
- Best-in-class PSRR
- Any format on any output: LVPECL, LVDS, HCSL, CMOS, HSTL, SSTL

LOW JITTER OSCILLATORS

- Wide frequency range: 10 MHz to 1.4 GHz
- Low jitter: 0.3 ps RMS
- Single, dual, quad & any frequency configurations
- Samples of any frequency available in 2 weeks
- LVPECL, LVDS, CML & CMOS options
- 1.8, 2.5 or 3.3 V options
- Industry-standard, RoHS-compliant 5 mm x 7 mm package and pinout

SILICON OSCILLATORS

- Quartz-free, all silicon construction
- Generates any frequency from 0.9 to 200 MHz
- 1.5 ps RMS phase jitter; 2 ps RMS period jitter
- LVPECL, LVDS, CMOS, HCSL & SSTL options
- 1.8, 2.5 or 3.3 V options
- Industry-standard 3.2 mm x 5 mm form factor

Customize an application-specific clock or oscillator today
at **www.silabs.com/custom-timing**

SILICON LABS

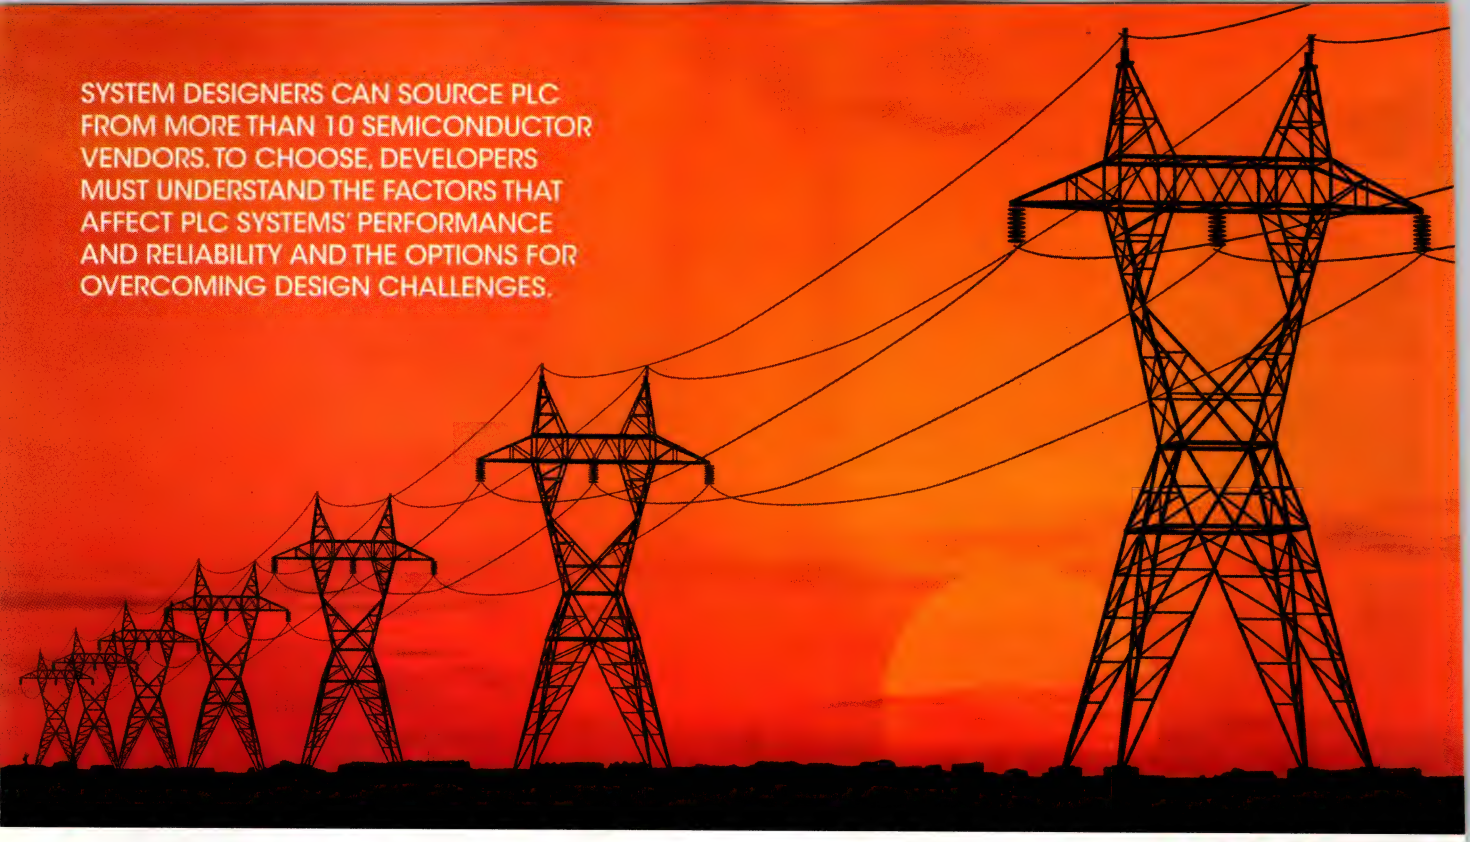

SYSTEM DESIGNERS CAN SOURCE PLC FROM MORE THAN 10 SEMICONDUCTOR VENDORS. TO CHOOSE, DEVELOPERS MUST UNDERSTAND THE FACTORS THAT AFFECT PLC SYSTEMS' PERFORMANCE AND RELIABILITY AND THE OPTIONS FOR OVERCOMING DESIGN CHALLENGES.

DESIGNING RELIABLE POWER-LINE COMMUNICATIONS

BY ASHISH GARG AND ANGAD SINGH GILL • CYPRESS SEMICONDUCTOR

PLC (power-line-communication) technology uses power lines as its transport medium. The data travels over the same power line that provides electricity, thus allowing the infrastructure in homes or cars to also transport data without adding wires. PLC technology is now experiencing rapid growth, finding its way into multiple applications and market segments, including the smart grid, lighting control, solar-panel monitoring, energy metering, in-home video distribution, and electric cars.

The global push for energy conservation is driving the need to intelligently communicate with both energy-generating and energy-consuming devices. PLC offers a unique no-new-infrastructure approach to enabling rapid deployment of smart energy-management technology worldwide. Unlike wireless approaches, PLC lacks the limitations of line-of-sight issues and short transmission range. PLC is also a cost-effective and easy-to-install technology for many applications.

Any communication system comprises a transmitter, a receiver, the medium, and a signal. In a generic PLC system, the transmitter modulates and injects the signal into the power line (**Figure 1**). The receiver at the opposite end of the link demodulates the signal and retrieves the data. The impedance of the power line attenuates the signal as it travels from the transmitter to the receiver. Any noise in the medium also corrupts the signal as it moves through the power line. The factors that affect the performance and reliability of a PLC system include the transmit-signal strength, the noise on the power line, the impedance of the power-line network, the protocol in use, and the receiver's sensitivity.

TRANSMIT SIGNAL

A stronger transmit signal implies more signal power across the line.

Support **Across The Board.**™
From Design to Delivery™

Now, you can have it all.™

Faster and easier than ever before. Our commitment to customer service is backed by an extensive product offering combined with our supply chain and design chain services – which can swiftly be tailored to meet your exact needs. We have dedicated employees who have the experience to provide the highest level of customer service with accuracy and efficiency. All of our technical experts are factory certified on the latest technologies, providing you the expertise to move projects forward with speed and confidence.

Avnet offers the best of both worlds: extensive product and supply chain knowledge, and specialized technical skill which translates into faster time to market – and the peace of mind that comes from working with the industry's best. **Avnet is ranked Best-In-Class* for well-informed sales reps, knowledgeable application engineers and our design engineering services – proof that we consistently deliver:**

> Industry recognized product expertise
> Specialized technical skills

Ready. Set. Go to Market.™
Visit the Avnet Design Resource Center™ at:
www.em.avnet.com/drc

SAME-DAY SHIPPING
orders placed before 8:00PM central time

AVNET® electronics marketing

Accelerating Your Success™

50 Years AVNET ON THE NEW YORK STOCK EXCHANGE NYSE.

Supplier Authorized Distributor

1 800 332 8638
www.em.avnet.com

Follow us on Twitter!
www.twitter.com/avnetdesignwire

*As rated by Hearst Electronics Group: The Engineer & Supplier Interface Study, 2009.

Stronger signals are less prone to the corrupting effects of noise on the power line and can travel farther. Transmit-signal strength also affects the PLC node's power consumption because the node consumes more power as more signal energy enters the line.

In the best-case scenario, developers would increase the signal strength of the transmitter until they achieved the best performance and power consumption over the power line. However, organizations such as the FCC (Federal Communications Commission) in North America and CENELEC in Europe tightly control transmit-signal strength. The FCC and CENELEC also regulate the harmonics that the main transmitter signal can inject into the power line. These regulations prevent signals on different frequency bands from corrupting one another.

When selecting a PLC device, check that it meets the transmit-signal-strength requirements for your target market. It should also comply with the standards that the FCC and CENELEC set. Ideally, the transmitting gain should be configurable so that you can tune signal strength based on the rest of the system. Additionally, confirm how much energy the PLC node consumes to achieve the best transmit-signal strength that the FCC and CENELEC require.

AT A GLANCE

- PLC (power-line-communication) transmit-signal strength is subject to target-market regulations and concerns about practical power consumption.
- Impulse and continuous noise derive from different sources and have different mitigation strategies.
- Varying and accumulated impedance attenuates the signal as it traverses a path from the transmitter to the receiver.
- Boosting the receiver sensitivity is not sufficient for improving communication-channel reliability; deft differentiation between the signal and its associated noise is also necessary.
- Highly integrated devices can positively affect end-system cost, and coexistence is also a valued attribute in today's fragmented and fast-evolving PLC market.

NOISE ON THE POWER LINE

Once the transmit signal has been injected into the power line, its integrity depends on the amount of noise on the line; stronger noise does greater damage to the signal. Noise can come from multiple sources. Simplistically speaking, noise on the line subdivides into impulse and continuous noise. Impulse noise is unpredictable and occurs in bursty sequences (**Figure 2**). This type of noise can come, for example, from a switched-on blender in the kitchen. It can be difficult to design a system that can tolerate the unpredictability and magnitude of impulse noise without compromising its data rate. This type of noise often obliterates any data packets on the line.

Continuous noise is more predictable than impulse noise (**Figure 3**). Continuous noise is usually a function of the

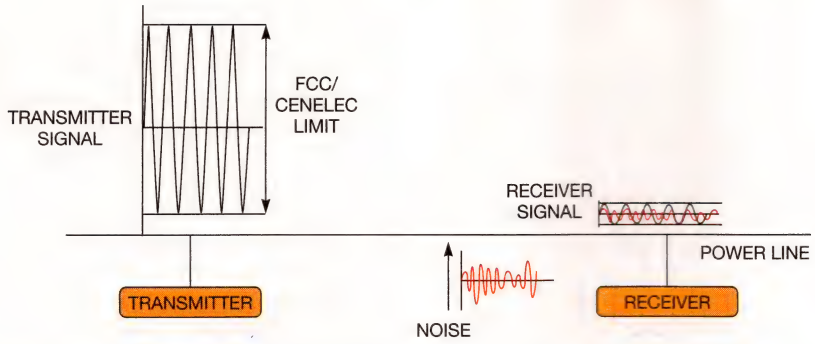

Figure 1 The impedance of the power line, along with noise on the line, can significantly attenuate the transmitted signal.

quality of the region's power-line installation. The developers of the power-line infrastructure designed it to efficiently carry power, not data, so they paid little attention to noise levels during power-line installation. Power lines' noise level depends on what part of the world a system operates in.

To enable robust communication over the power line, the SNR (signal-to-noise ratio) must remain above a certain threshold. If high-amplitude, continuous noise exists within the frequency range of the PLC system, you should isolate the noise by moving it away from PLC receivers or by adding a blocking inductor to the power supply of the noise-generating equipment, thereby attenuating the noise frequency below the receiver's SNR.

Developers can also use several other techniques to overcome the effects of noise. These techniques include the use of bidirectional communication, retries, error detection, and AGC (automatic gain control). If a PLC system communicates in only one direction, the transmitter cannot know whether communication succeeds. This shortcoming was one of the biggest of the original unidirectional X10 PLC technology. Bidirectional communication allows the receiver to send an acknowledgment after successful reception of data. In case the receiver does not receive an acknowledgment, the transmitter can take corrective action.

In a bidirectional system, communication confirmation can occur through an acknowledgment mechanism. If an intelligent transmitter does not receive a reply from the receiver, then it can resend data packets. A PLC's implementation of built-in retries can be a powerful means of achieving reliable PLC. Even after a receiver successfully receives a data packet, it still must check it for any noise-caused damage. CRCs (cyclic redundancy checks) enable the receiver to detect any data-packet errors. When the receiver detects an erroneous data packet, it can either request the transmitter to resend it or not acknowledge the data, triggering

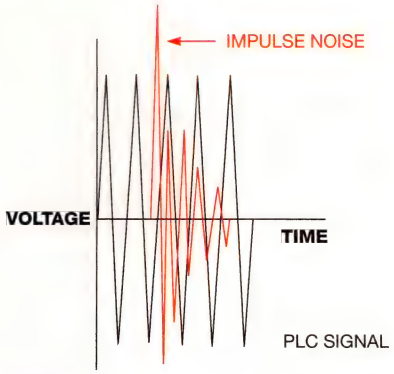

Figure 2 Impulse noise on the power line is unpredictable and can obliterate any coincident data packets.

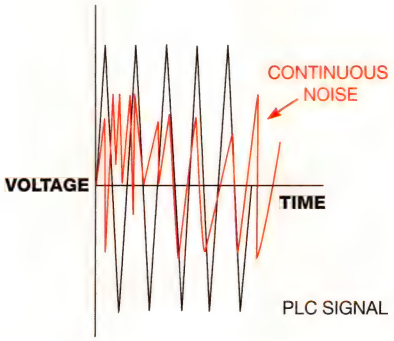

Figure 3 You can more easily predict and, therefore, compensate for continuous noise on the power line.

the transmitter to automatically retry transmitting a data packet.

To overcome the effects of continuous noise, some PLC devices implement AGC. Using AGC, the receiver dynamically adjusts its sensitivity above the noise floor so that it can better differentiate between noise and data. Clearly, the more ways a system can accommodate or overcome noise, the more reliable the system. Hence, it is beneficial to implement acknowledgment-based bidirectional communication with retries and CRC.

IMPEDANCE

The impedance that a signal sees on a power line affects the signal power that the transmitter can transfer into the power line. This impedance depends on the impedance of the power line and that of the nodes—that is, the appliances—that connect to it. Power-line impedance changes whenever you plug an appliance or a node into a power socket. Maximum signal power transfers when the impedance that the signal sees in the power line matches that of the transmitter circuit. The greater the difference between these two impedances, the less the transferred signal power; as a result, PLC performance degrades. The dynamic change in impedance over time is one of the toughest issues to address in power-line applications. PLC transmitters and receivers must anticipate these impedance changes in the power line if they are to achieve robust signal performance. Continually matching the impedance of the transmitter to that of the power line allows maximum signal transfer, and high receiver impedance ensures minimal signal loss on the receiver side.

NETWORK PROTOCOL

A robust and error-free network protocol has perhaps the greatest impact on the reliability of PLC. Although the system design has little control

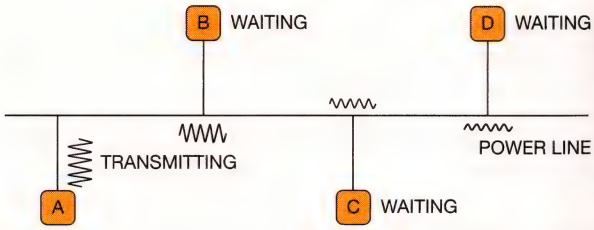

Figure 4 The CENELEC-mandated CSMA scheme ensures that multiple nodes can coexist on a power line, efficiently sharing access.

over physical factors, such as noise and power-line impedance, a power-line-optimized network-protocol implementation can significantly improve PLC performance. Your network-protocol selection can make or break a PLC system; with the right network protocol, you can achieve 100%-successful PLC communication. Consider that most PLC applications support tens to hundreds of nodes that connect on the same power line. The network protocol arbitrates data packets between nodes so that all nodes can fairly share the available bandwidth on the line and no one node can monopolize the communication channel. The definition and implementation of the network protocol also determine the maximum number of PLC nodes that can communicate over a line.

The network protocol can comprehend acknowledgement, retries, and CRC. The application running on the PLC system, in such cases, need not worry about implementing these techniques. From an application standpoint, the software receives only valid PLC data. Some PLC devices come with a built-in network protocol, whereas others require developers to define, code, and manage their own protocols. If the protocol cannot run on the PLC device itself, the developer must specify another processor to implement the protocol.

Another important aspect of protocol selection is interoperability and coexistence. The CENELEC-mandated CSMA (carrier-sense-multiple-access) implementation ensures that one set of PLC nodes can coexist with those of other vendors (**Figure 4**). With the growth in PLC-enabled devices, coexistence is an important way to future-proof the deployment of PLC devices.

RECEIVER SENSITIVITY

Depending on the power line's characteristics, the loads, and the segment lengths that signals must travel over the power line, signals can be significantly attenuated by the time the receiver picks them up. A receiver that has a

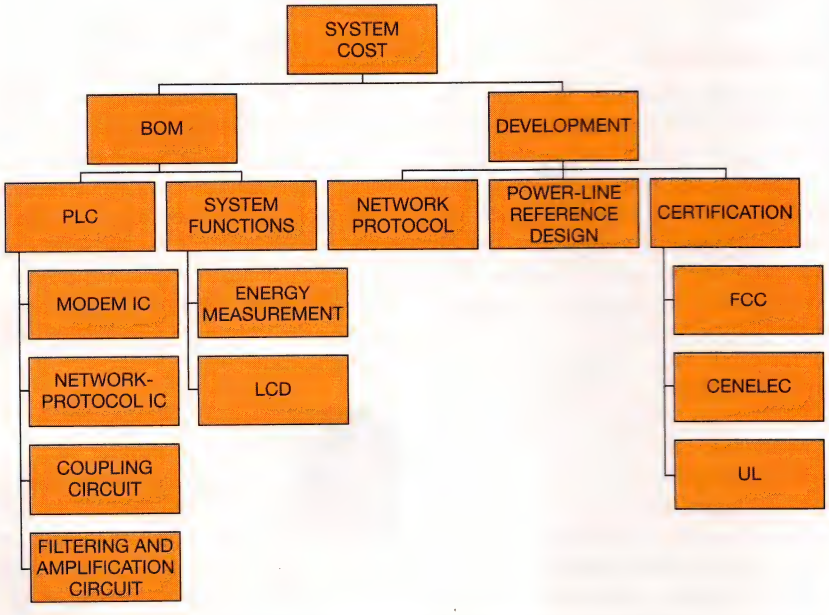

Figure 5 The costs of implementing PLC group into BOM and development categories.

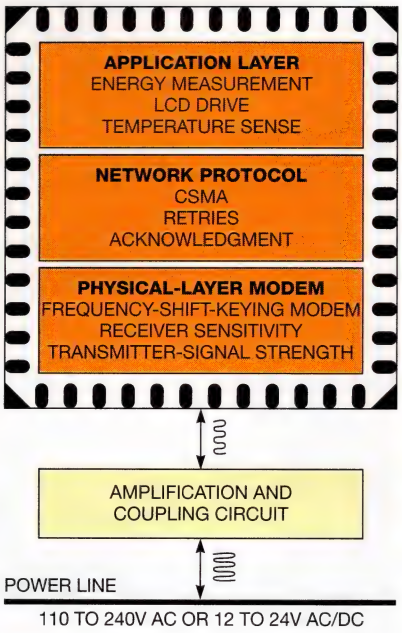

Figure 6 The Cypress CY8CPLC20 PLC device integrates a physical-layer modem, a network protocol, and an application layer.

high receiver sensitivity—that is, it can reliably receive low signal strengths—can pick up even lower-strength signals from the line, thereby increasing the effective communication distance. However, high sensitivity is not always a good attribute. A high-sensitivity receiver, for example, detects not only small signals but also small amounts of noise in the channel. Therefore, it is important to have a mechanism in place that prevents the receiver from confusing noise as the actual signal. AGC is one such mechanism. The receiver uses AGC to dynamically adjust its sensitivity above the noise floor so that it can better differentiate between noise and data.

MULTIPLE PHASES

Most buildings contain multiple ac phases. Transformers operating at 50 or 60 Hz produce these phases. Most PLC signals operate at higher frequencies, however, so the transformer may filter out the signal and consequently cannot jump to the adjacent phase, which may be in the same building. This scenario creates a potential problem in which the PLC signal cannot reach all sockets in a house or another structure. However, this outcome depends on the design of the transformer. You can solve this problem by coupling the PLC signal from one phase to another. Well-known techniques for accomplishing this objective include capacitive phase coupling and wireless phase coupling.

Capacitive phase coupling requires the connection of a capacitor across the phases at the transformer, thereby allowing the PLC signal to pass. It also requires physical access to the transformer; hence, this approach might be unfeasible or not cost-effective in some circumstances. In wireless phase coupling, the PLC data transmits from one phase to another using two RF (radio-frequency) devices, one of which connects to each phase. These two devices can connect to any socket on the phases, as long as they are in range of each other. This technique requires no physical access to the transformer.

Wireless coupling is less intrusive to the transformer; thus, many designers prefer it over capacitive coupling. Some PLC devices come with wireless-coupling options, whereas others leave it to designers to develop a way to couple phases.

PLC ARCHITECTURES MUST MINIMIZE SYSTEM COST TO EFFECTIVELY COMPETE IN THE MARKET.

SYSTEM COST

Although reliability is a key design factor, PLC architectures must also minimize system cost to effectively compete in the market. When adding PLC to a system, some designers aggressively negotiate chip price but ignore the overall cost of adding PLC functions to their systems, warranting a more holistic look at PLC cost.

You can broadly categorize PLC costs into BOM (bill-of-material) and development costs (**Figure 5**). The BOM cost includes the cost for all the ICs and components that make up the system, including PLC and other system-relevant functions. Development costs, on the other hand, include the costs of other resources, such as the network-protocol implementation, board and layout design, and expenses for certifying the product with FCC, CENELEC, and UL (Underwriters Laboratories) standards.

Developers can reduce overall system cost by integrating as many of these costs into the fewest devices. For example, a device that provides modem- and network-protocol support on one IC is less expensive than the alternative approach of implementing the functions using separate ICs. Single chips integrating system functions such as energy measurement, LCD driving, temperature sensing, and load control also lower system cost and reduce integration and development complexity (**Figure 6**). Many PLC controllers also have reference designs complying with FCC, CENELEC, and UL standards to further accelerate development. In general, the more complete a PLC device is, the less time and money it takes to implement PLC. EDN

AUTHORS' BIOGRAPHIES

Ashish Garg is a product-marketing manager at Cypress Semiconductor, where he manages the dc/dc- and PLC-product lines. Before joining Cypress, Garg worked in various marketing and technical positions at Altera, Texas Instruments, and Sun Microsystems. He has a master's degree in business administration from the Wharton School of the University of Pennsylvania (Philadelphia), a master's degree in electrical engineering from the University of California—Santa Barbara, and a bachelor's degree in electronics engineering from the Birla Institute of Technology and Science—Pilani (India).

Angad Singh Gill is a product-marketing engineer at Cypress Semiconductor, where he is responsible for the dc/dc-power- and PLC-product lines. Gill has a master's degree in physics and a bachelor's degree in electronics engineering from the Birla Institute of Technology and Science—Pilani (India).

SPARTAN®-6

- Easy-to-use, low power, serial transceivers support up to 3.125Gbps to enable industry standards such as PCIe®
- Low voltage option reduces total power consumption by 65% over previous generations
- Integrated DSP, memory controllers, and clocking technology simplifies designs

VIRTEX®-6

- High bandwidth serial connectivity with up to 72 low-power transceivers supporting up to 11.18Gbps
- Ultra high-performance DSP using up to 2016 low-power, performance-optimized DSP slices
- Integrated high-performance ExpressFabric™ technology running at 600 MHz clocking and performance-tuned IP blocks
- Proven cost-reduction with EasyPath™-6 FPGAs

Potential. Realized.

Unleash the full potential of your product design with Xilinx® Virtex®-6 and Spartan®-6 FPGA families — the programmable foundation for Targeted Design Platforms.

- Reduce system costs by up to 60%
- Lower power by 65%
- Shrink development time by 50%

Realize your potential. Visit www.xilinx.com/6.

XILINX®

USE GPUs TO BOOST ACCELERATION

BY GRAHAM PROPHET • EDITOR, *EDN EUROPE*

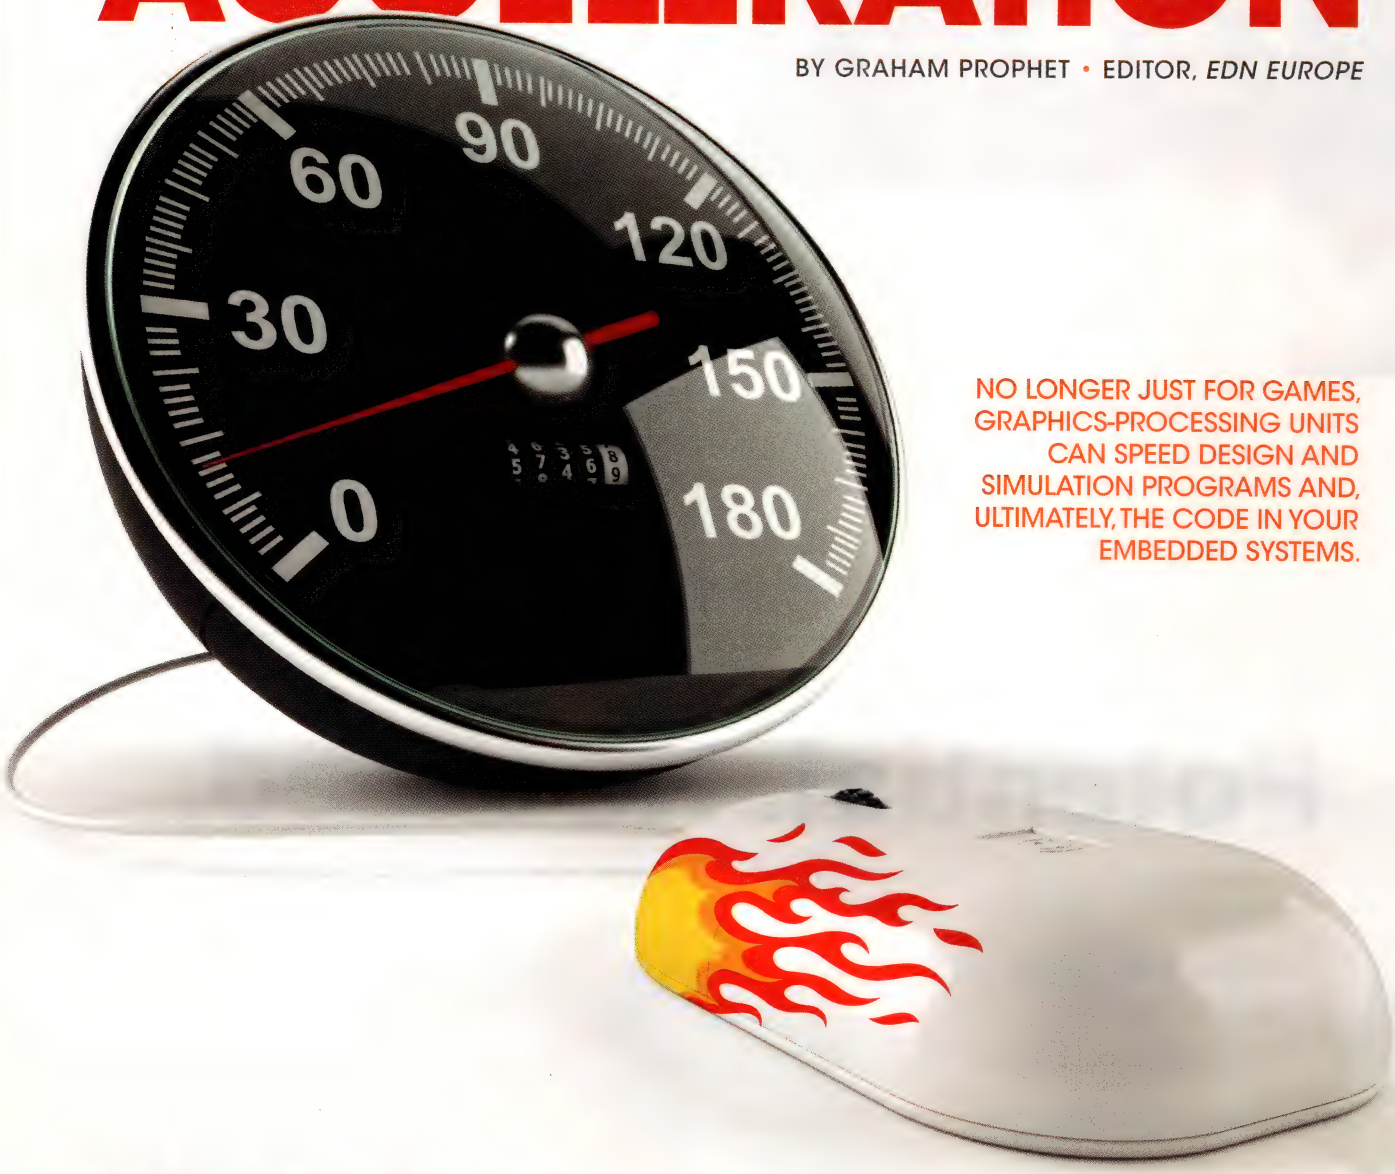

NO LONGER JUST FOR GAMES, GRAPHICS-PROCESSING UNITS CAN SPEED DESIGN AND SIMULATION PROGRAMS AND, ULTIMATELY, THE CODE IN YOUR EMBEDDED SYSTEMS.

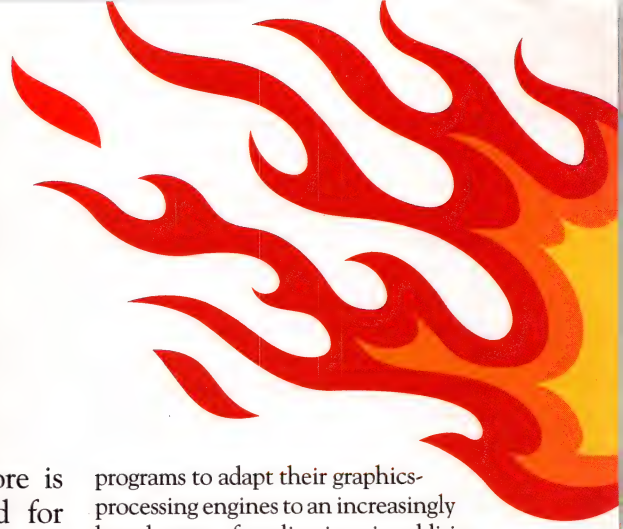

In desktop, laptop, and notebook computing, multicore is now mainstream; single-core processors that pushed for higher and higher clock speeds long ago reached the point of diminishing returns, and today's baseline PC now typically sports not only a quad-core CPU but also a multicore GPU (graphics-processing unit). The same processors or, rather, the industrial long-term-support versions of them, are similarly well-established in single-board computing and other embedded-computing platforms.

PCs with even modest specifications also contain a processing unit with tens or hundreds of cores. That chip is a GPU (graphics-processing unit), typically from AMD (Advanced Micro Devices) or Nvidia. In the United Kingdom, approximately £100, or less than $150, will buy you, for example, a 192-core PCIe (Peripheral Component Interconnect Express) graphics card that boasts a processor clock speed of 1566 MHz and a memory bandwidth of 57 Gbytes/sec to its 1 Gbyte of graphics RAM. The GeForce GTS 450 is a midrange unit by the standards of enthusiastic gamers, who will happily extend their graphics-card budget to the $500 to $800 range. That amount of cash will buy cards such as the ATI Radeon HD 5970, with dual GPUs and no fewer than 3200 processing cores, 2 Gbytes of graphics RAM, and all the high-end polygon-drawing specifications that purchasers in that sector of the market expect.

For a few hundred dollars, therefore, you can put teraflops-class computing capability into a PCIe slot. If your primary interest doesn't lie in painting the screen at the highest possible pixel rate, can you use that capability for other purposes—for example, to accelerate the programs you use to design and simulate embedded systems? You can, and the tools to help you do so are becoming more and more widespread. The generic name for the approach is GPGPU (general-purpose computing using GPUs).

CONCEIVED FOR GRAPHICS

An immediate caveat is that the individual processing units in a GPU are expressly for tasks such as calculating and rendering polygons, texture mapping, and calculating pixel values and pushing them to a screen at the fastest possible rate. Adapting a GPU to your needs in other directions is most likely to reward your efforts if you have a task that has a high degree of inherent parallelism and that maps well onto the GPU architecture.

More recent GPU designs, however, allocate more silicon features for broadening the scope of the tasks you might allocate to them. Current-generation GPUs implement ECC (error-correction code), for example. One memory-read or -write error—a pixel or a polygon momentarily the wrong color, perhaps—when the chips output only graphics is of little consequence. That situation is not helpful for high-performance computing, however, and manufacturers quickly remedied it.

GPGPU uses the GPU as a form of stream processor, with a model similar to SIMD (single instruction/multiple data), in which many items of data undergo the same processing operation; both AMD and Nvidia have ongoing programs to adapt their graphics-processing engines to an increasingly broad range of applications in addition to high-performance graphics. In the specifications for their graphics cards, look for the terms ATI Stream processing and Nvidia CUDA (compute-unified device architecture). Nvidia now hardly ever expands the acronym CUDA, using it as a brand in its own right.

Presenting a keynote address at Nvidia's GPU Technology Conference in September in San Jose, CA, Nvidia's president, Jen-Hsun Huang, dated the transition to four years ago. "We realized that the 3-D graphics processor could evolve into something more general. ... We could combine the multicore CPU with the many-core GPU for parallel processing. ... There are many applications where 5% of the code takes up the vast majority of the runtime [in which such an architecture would be useful]," he said. "We made the conscious decision to ride the back of the high-volume [graphics-card] business and to take advantage of the $1 billion-a-year R&D budget for the GeForce line ... to make high-performance computing more affordable."

Both AMD for ATI Stream and Nvidia for CUDA provide SDKs (software-development kits) to program their silicon. Both provide what amount to C-language extensions to provide the means for coders to format suitable segments of their overall tasks, pass those jobs to the GPU, execute the operations at which the GPU is most adept, and return the results of the GPU's deliberations to the main program flow.

Both companies also support the OpenCL (Open Computing Language) standard, which the Khronos Group administers and which aims to provide a means of accelerating parallel code that is portable across multiple processor architectures—both GPUs and other

CPUs with and without acceleration. OpenCL's developers named it to be analogous with OpenGL (Open Graphics Language), the widely used graphics-programming standard, in acknowledgment of the languages' common use of the GPU platform, but the nomenclature may be more confusing than helpful. In June 2010, Khronos, in its own words "an industry consortium creating open standards to enable the authoring and acceleration of parallel computing, graphics, and dynamic media on a wide variety of platforms and devices," released the OpenCL 1.1 specification, which includes a C++ wrapper API (application-programming interface) for use with the standard.

"The OpenCL 1.1 specification is being released 18 months after OpenCL 1.0 to enable programmers to take even more effective advantage of parallel-computing resources while protecting their existing investment in OpenCL code," says Neil Trevett, chairman of the OpenCL working group, president of the Khronos Group, and vice president of Nvidia.

OpenCL 1.1, the organization says, adds more functions to increase parallel-programming flexibility and performance, including, for example, data types such as three-component vectors and additional image formats. You can gain an indication of the pace of proliferation of the GPGPU concept from the fact that the OpenCL 1.1 working group includes names such as Apple, which originated the OpenCL project; ARM; Broadcom; Ericsson; Freescale; IBM; Intel; Nokia; Renesas; ST-Ericsson; STMicroelectronics; Symbian; and Texas Instruments. According to Pete Hutton, vice president of technology and systems at ARM's processor division, the company is committed to supporting OpenCL across a range of its products. "OpenCL 1.1 will accelerate innovation in high-performance-computing systems containing GPUs, CPUs, and other devices [from] ARM partners," he says.

"Has OpenCL advanced sufficiently that I don't have to learn [the native format of the GPU maker]?" asks a contributor to one of the user forums associated with GPGPU. He receives this—slightly cautious—reply from a fellow user: "I think so: There are bugs, but there are bugs with any new software."

AMD's Web site provides tutorial material from its perspective on working in OpenCL, describing how its objective is to accelerate parallel code that will run across a range of both CPUs and GPUs. OpenCL, explains a video that is part of the tutorial content, supports a kernel-based model, which can be either data-parallel or task-parallel. A host machine connects to one or more OpenCL devices, which may be in SIMD or SPMD (single-program/multiple-data) fashion. The user acquires, with OpenCL, the means to queue work for submission to the GPU. The AMD speaker remarks on the need for a rigorous approach to memory management when coding for OpenCL. OpenCL recognizes several classes of memory, including private, local, local-global, and host, for which there are no overall synchronization mechanisms. The user must manage data flow to and from memory and explicitly handle synchronization.

Whether under OpenCL, CUDA, or ATI Stream, you have the advantage that your development platform investment is small: just a PC with a suitably enabled graphics card. Alternatively, under plans that PGI (Portland Group Inc) recently announced, you may not require even that approach. The high-performance-computing-compiler company, a subsidiary of STMicroelectronics, has been developing a CUDA C compiler targeting 64- and 32-bit x86 architectures that it planned to demonstrate last month at the SC10 Supercomputing conference, which took place in New Orleans. The compiler allows you to compile and optimize CUDA applications to run on x86-based workstations, servers, and clusters with or without an Nvidia GPU. If your host lacks a GPU, applications use multiple cores and the streaming-SIMD capabilities of Intel and AMD CPUs for parallel execution. The company has developed the code in response to developers who want to use a single parallel programming model to target many-core GPUs and multicore CPUs.

AT A GLANCE

- For a few hundred dollars, you can put teraflops-class computing capability into a PCIe (Peripheral Component Interconnect Express) slot.
- Both AMD (Advanced Micro Devices) and Nvidia have ongoing programs to adapt their graphics-processing engines to an increasingly broad range of applications in addition to high-performance graphics.
- Whether under OpenCL (Open Computing Language), CUDA (compute-unified device architecture), or ATI Stream, you have the advantage of a small development-platform investment.
- Matlab represents vectors and matrices in an array format that aligns readily with the array-oriented operation of the GPU (graphics-processing unit).
- Ansys has demonstrated GPU acceleration on both mechanical- and electrical-simulation tools.
- Agilent Technologies' EEsof division and Nvidia have developed a relationship with symmetry.

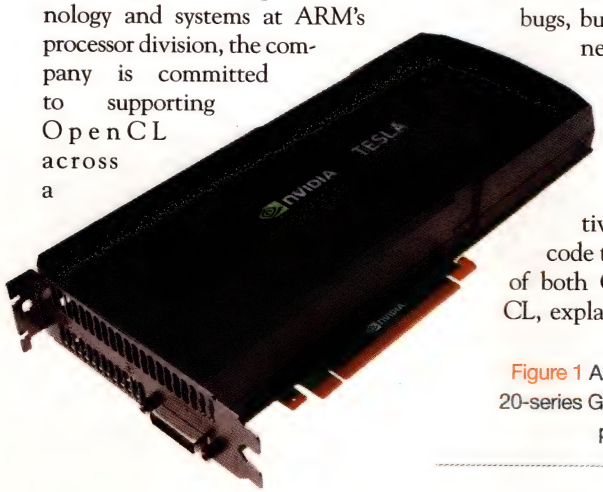

Figure 1 Ansys reports that Nvidia Tesla 20-series GPUs, such as the C2050, can produce faster solution times.

PARALLEL COMPUTING

An illustration of where GPGPU might take you comes from The MathWorks, which has announced support for Nvidia GPUs for its simulation products and has recoded a menu of functions to take advantage of Nvidia GPUs (**Listing 1**). The user does not see the detail of the GPU coding but can use control statements to allocate parts of the program flow to the GPU. In this signal-processing example, the panel on the left depicts the code to generate

40% bigger screen.

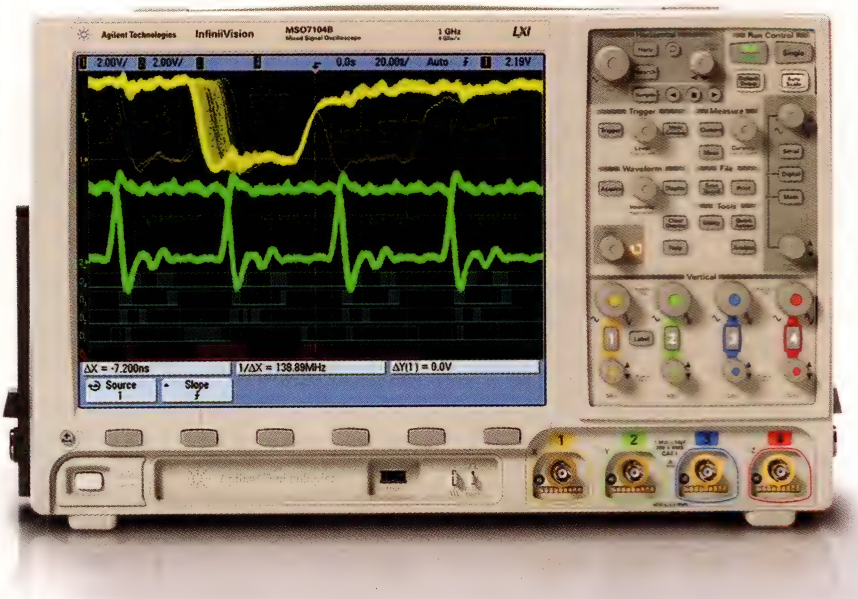

475,000 times faster...Yes, really.

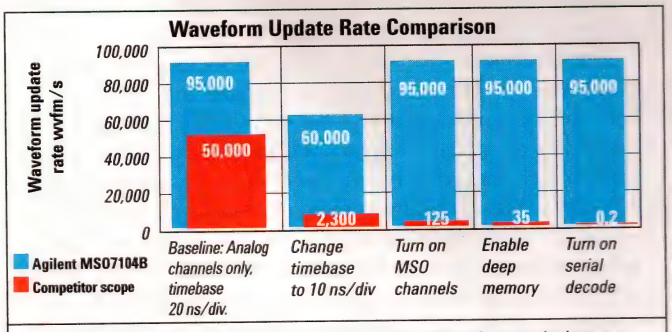

Notes:
- Measurements taken on the signal using Agilent MSO7104B and a competitor's scope.
- Memory depth = display window X sample rate with up to 8 Mpts on Agilent.
- Competitor scope's test measurements taken with 2.13 firmware.

A scope's real world performance depends on more than just bandwidth and sample rates. It takes an ultra fast update rate to capture infrequent events and display the most accurate signal detail. That's where Agilent 7000B Series scopes with bandwidths from 100 MHz–1 GHz really stand out. Thanks to a proprietary chipset, you get faster update rates not only in analog mode, but in every mode. That's why you see greater signal detail. That's Agilent.

Agilent and our Distributor Network

Right Instrument.
Right Expertise.
Delivered Right Now.

Newark™

800-463-9275
www.newark.com/agilent

Are you using the best scope?
Take the 5-minute Scope Challenge.
www.newark.com/Agilent_Scopes

Agilent Technologies

LISTING 1 SIGNAL-PROCESSING EXAMPLE

```
D = data;
iterations = 2000; % # of parallel iterations
stride = iterations*step; %stride of outer loop

M = ceil((numel(x)-W)/stride);%iterations needed
o = cell(M, 1); % preallocate output

for i = 1:M
   % What are the start points
   thisSP = (i-1)*stride:step: ...
      (min(numel(x)-W, i*stride)-1);

   % Move the data efficiently into a matrix
   X = copyAndWindowInput(D, window, thisSP);

   % Take lots of fft's down the columns
   X = abs(fft(X));

   % Return only the first part to MATLAB
   o{i} = X(1:E, 1:ratio:end);

end
```

```
D = gpuArray(data);
iterations = 2000; % # of parallel iterations
stride = iterations*step; %stride of outer loop

M = ceil((numel(x)-W)/stride);%iterations needed
o = cell(M, 1); % preallocate output

for i = 1:M
   % What are the start points
   thisSP = (i-1)*stride:step: ...
      (min(numel(D)-W, i*stride)-1);

   % Move the data efficiently into a matrix
   X = copyAndWindowInput(D, window, thisSP);

   % Take lots of fft's down the columns
   X = gather(abs(fft(X)));

   % Return only the first part to MATLAB
   o{i} = X(1:E, 1:ratio:end);

end
```

a spectrogram using the CPU alone; highlighted keywords in the center and right panels' code segments show the changes that place the array calculations, first on a single GPU and then on parallel GPUs.

Jos Martin, head of the company's parallel-computing-development team, explains The MathWorks' interest. Both ATI and Nvidia, Martin notes, anticipated the general user base's interest in using their chips for more general computing and added silicon support to their products that supported a wider range of functions than rendering. Some of the added features included support for double-precision floating-point arithmetic, linear algebra, and more matrix operations. For Matlab, Martin's group focused on the Nvidia offering, which includes the IEEE-compliant double-precision floating-point capability. Given that feature and the ability to efficiently compute functions such as FFTs (fast Fourier transforms) in the GPU, it became worthwhile to build a set of numerical libraries.

Matlab, Martin says, has vectors and matrices as its basic number representations in the array format that aligns readily with the array-oriented operation of the GPU. In its graphics-derived mode of operation, the GPU performs the same operation on each data point in an array. Each operation constitutes a kernel that executes in its own memory space and with no communication between kernels. Matlab's data format naturally parallelizes into this environment; the addition of linear algebra provides functions such as matrix multiplication. Data can reside in the GPU memory for the duration of each operation that you delegate to the GPU. The Matlab team has now implemented approximately 100 functions in its parallel-computing tool kit, each using the same method: You load the relevant data into an array, "push" the data out to the GPU, specify the operation to be performed, and then "gather" the resulting data array back to the CPU. Overall program execution is, typically, four to 10 times faster than on a CPU, Martin explains; FFTs, in particular, run much faster due to the high bandwidth between memory and each GPU.

The Matlab team hides the CUDA structure of threads and blocks. This approach, Martin says, can "gloss over how hard the programming model is." For example, he adds, "There can

```
D = gpuArray(data);
iterations = 2000; % # of parallel iterations
stride = iterations*step; %stride of outer loop

M = ceil((numel(x)-W)/stride);%iterations needed
o = cell(M, 1); % preallocate output

parfor i = 1:M
   % What are the start points
   thisSP = (i-1)*stride:step: ...
      (min(numel(D)-W, i*stride)-1);

   % Move the data efficiently into a matrix
   X = copyAndWindowInput(D, window, thisSP);

   % Take lots of fft's down the columns
   X = gather(abs(fft(X)));

   % Return only the first part to MATLAB
   o{i} = X(1:E, 1:ratio:end);

end
```

be many blocks of processing threads that cannot talk to each other, whereas other entities can communicate over shared memory. There is no global-memory synchronization. You have to be very careful how you access memory, in a context in which the aggregate memory bandwidth can be 140 Gbytes per second."

SPEEDING SIMULATIONS

Announcements of support for GPUs in application programming have also come from Ansys and Agilent Technologies' EEsof division. Ansys announced its HPC (high-performance-computing) approach, which it based on Nvidia CUDA-based Tesla GPUs (**Figure 1**). According to Ansys, HPC can cut overall engineering-simulation time on a quad-core CPU and Tesla GPU by as much as half compared with the quad-core CPU alone. Ansys has demonstrated the capability on both mechanical- and electrical-simulation tools in a preview release and plans a full Ansys accelerator capability for general availability by year-end.

The overall reduction for a complete simulation implies that the portion that the GPU accelerates achieves a higher speedup. The company obtained that figure in its Ansys Mechanical product but says that the benefits for other simulations that use similar mathematical techniques, such as finite-element analysis—for example, 3-D-EM (electromagnetic)-field software—should show gains at least as high (**Figure 2**). The GPU accelerates only the numerical-analysis portion of the simulation. A paper by Ansys developers outlines trials in which the host machine contained a four-core Intel Xeon 5560, and the GPU card used a Tesla 20-series C2050. Ansys is also investigating the use of AMD ATI GPU cards in a similar role.

The relationship that has developed between Agilent Technologies' EEsof division and Nvidia has symmetry; Agilent announced about three years ago that it was working on GPU support of its simulation tools. Nvidia uses the same tools for signal-integrity simulations in its own IC-design flow. Agilent offers acceleration of elements of its circuit, device, and EM-field tools using GPUs, again for the CUDA architecture. In the frequency domain, Agilent has applied the technology to its Spice-based simulations of elements such as backplanes, long traces, and connectors in high-frequency digital systems. Time-domain simulations are more amenable to a parallel-computation approach, the company says. Finite-difference analysis of large—relative to wavelength dimensions—physical spaces provides a suitable task—for example, the antenna system in a mobile handset, or radiated fields in a vehicle. An application that has demonstrated similarly good results is analysis of EMI (EM interference) from a high-speed digital PCB (printed-circuit board). This analysis requires a fine-mesh model with small cells and, using the GPU, shows as much as an order-of-magnitude speedup in simulation.

The Spice problem, according to Colin Warwick, Agilent's EEsof signal-integrity-product manager, is inherent-

ly parallel and maps readily onto the GPU array. Speaking of simulations of as many as 90,000 transistors, Warwick says, "80% of a Spice run [the numeric solution] can go down to nearly 0%." When you take into account the parts of the program that are more difficult to parallelize and that must still run on the CPU, overall fivefold decreases in runtime are possible. These gains become large when you are doing Spice runs of millions of bits passing through a digital channel to build an eye diagram.

Agilent has, like other companies, written the relevant routines for the CUDA architecture. And, as with Ansys, Agilent's base machine has a four-core processor and at least 8 or 16 Gbytes of memory. Adequate memory is key to preserving the gains of the GPU when farming out large parallel tasks to the GPU. Early GPU products were limited in this respect, Warwick says, but he notes that current Nvidia cards using its Fermi architecture can host 6 Gbytes. A PCIe card cage with four slots yields 24 Gbytes of available memory. Nevertheless, the host still sees it as a standard PCIe device, so you have all the benefits of low-cost, widely deployed hardware.

Agilent also uses Nvidia GPUs in the role that their designers first envisaged: the display of high-resolution graphics, including 3-D visualization. A stereo-3-D-viewer capability

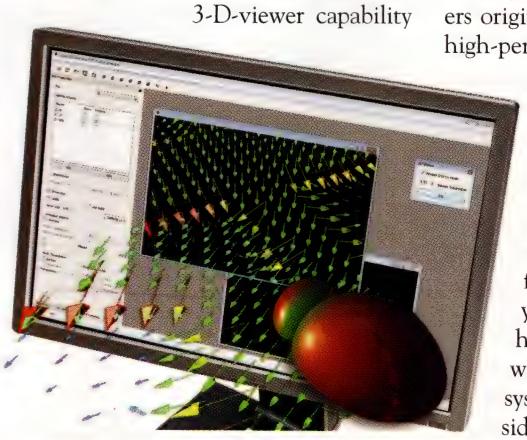

Figure 3 Arrows, representing current direction and magnitude at each point of a conducting structure, and hemispherical lobes, representing the far-field radiation patterns, seem to jump out of the screen with the ADS stereo viewer's 3-D effect. ADS uses Nvidia GPUs for graphics acceleration.

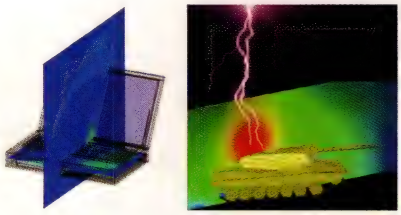

Figure 2 Ansys 13.0, which provides support for GPUs, includes a new electromagnetic transient solver that enables engineers to study the electromagnetic effects in broadband and radar applications as they occur over time.

exists for the Momentum G2 Element and FEM (finite-element-method) EM simulators in the ADS (Advanced Design System) EDA platform (**Figure 3**). The stereo-3-D viewer uses Nvidia Quadro GPUs and 3-D Vision quad-buffered stereo technology to render electric and magnetic fields and currents in high-resolution 3-D. You need a host machine with an Nvidia Quadro FX; Nvidia 3-D Vision stereoscopic glasses; and a 3-D Vision-ready, 120-Hz display. Viewing the simulation in 3-D reveals details and structure in the computed fields that you might easily miss when you view a flat screen, Warwick says.

STEPPING DOWN IN POWER

Accessing the resources of GPGPU through conventional PCIe cards is not a low-power approach. Their developers originally designed them to deliver high-performance graphics. The earlier-noted GTS 450 uses a maximum of about 100W and, like all other high-performance graphics cards, has its own cooling arrangements. If your interests lie in building ultimate-performance PC configurations for gaming, though, you know that more power-hungry products also exist, some with heat-pipe or fluid-cooling systems. Some may deliver considerable processing power but draw less power than their packaging implies. An element of marketing hype appears to reaffirm to the retail customer that the product must deliver impressive graphics because it sports a serious fan. Adding some macho-looking graphics or a bright-red fan to drive home the message is one tactic. As The MathWorks' Martin points out, you

should base your analysis on the power per floating-point operations per second in your complete system.

According to Nvidia's Huang, power will be the key constraint for future systems. "Performance per watt is the same thing as performance," he says. Huang sets out a road map for Nvidia's product line. He rates the current architecture, Fermi, successor to the Tesla-generation products, at about 1.5 double-precision Gflops/W. Using the same metric, he anticipates Nvidia's Kepler architecture, which uses 28-nm technology and which will enter production in 2011, at perhaps 4 double-precision Gflops/W, to be followed in 2013 by the Maxwell architecture at 16 double-precision Gflops/W.

If the company fulfills these plans, high-performance-computing developers should have resources to work with. Perhaps more relevant to many embedded-system designers looking to add a measure of GPGPU to their projects are the plans from both AMD and Nvidia for low-power product lines. Nvidia sells the graphics/system-chip Ion product family, which first targeted use in netbooks, into single-board-computer and embedded-computer manufacturers. Ion-family chips have as many as 16 CUDA-capable cores. The company is also rolling out its Tegra series of mobile processors; it aims Tegra 2 ICs, which have twin ARM9 cores, at the automotive sector, whereas the Tegra APX, with an ARM11 MPCore, targets the handheld-media-device market. Both devices contain an ultra-low-power, multiple-core Nvidia GPU.

ENTER THE APU

Meanwhile, AMD is releasing information on its first APU (accelerated processing unit), which it code names Llano and plans as the foundation of the Fusion family. AMD has scheduled formal introduction for early 2011. The Fusion series will combine x86 CPUs and a GPU that presumably will build on AMD's Radeon experience. AMD has confirmed that Fusion devices will serve computing devices at every level of power, including low-power and portable units, and that you will be able to access their GPU capabilities with the ATI Stream SDK or OpenCL.

As part of the OpenCL 1.1 release, Manju Hegde, AMD's corporate vice president for the Fusion Experience Program, remarks on the rapidly growing interest in GPU computing across the industry. "Vendors [must] embrace a multivendor, multisource interface and an industry-standard programming model," he says. "Fusion [will] unleash the true potential of application acceleration … with APUs designed to support OpenCL on both the GPU and the CPU, thereby providing a heterogeneous computing platform."

EMBEDDED GPGPU

EDA-tool providers such as Ansys and Agilent/EEsof bring the GPGPU phenomenon close to home for electronic-system designers. The technology appears ready, therefore, for you to use the approach to accelerate performance in a deeply embedded application. Today, if you Google "GPGPU" and "embedded," you'll mostly find references to what you might term "heavyweight" applications: radar, image processing, medical imaging, UAVs (unmanned aerial vehicles), and so on. But the tools and the hardware are in place to enable the power of these many-core devices for wider use. In many cases, the silicon is in place on your embedded single-board computer or in your desktop PC, just waiting for software that knows how to exploit multicore GPUs to best advantage. EDN

You can reach EDN Europe *Editor Graham Prophet at gprophet@reedbusiness.fr.*

FOR MORE INFORMATION

Agilent Technologies
www.agilent.com

AMD
www.amd.com

Ansys
www.ansys.com

Apple Inc
www.apple.com

ARM
www.arm.com

Broadcom
www.broadcom.com

Ericsson
www.ericsson.com

Freescale
www.freescale.com

IBM
www.ibm.com

Intel
www.intel.com

Khronos Group
www.khronos.org/opencl/

The MathWorks
www.mathworks.com

Nokia
www.nokia.com

Nvidia
www.nvidia.com

OpenCL
www.khronos.org/opencl

OpenGL
www.opengl.org

Portland Group
www.pgroup.com

Renesas
www.renesas.com

ST-Ericsson
www.stericsson.com

STMicroelectronics
www.st.com

Symbian
www.symbian.org

Texas Instruments
www.ti.com

BY JEFF GALLOWAY AND RANDY CAPLAN • SILICON CREATIONS

SOC-PLL design requires trade-offs

YOUR DECISION TO USE AN LC- OR A RING-BASED PLL DEPENDS ON THE APPLICATION.

PLLs (phase-locked loops) are common analog circuits in SOCs (systems on chips). Almost all SOCs with a clock rate greater than 30 MHz use a PLL for frequency synthesis. However, a "one-size-fits-all" PLL does not exist. The devices have a range of frequency, power, area, performance, and functions. PLLs implemented in 100 nm or smaller processes typically range in frequency from 10 MHz to 10 GHz. Their power spans from less than 1 mW to more than 100 mW. Their size can vary from 0.04 to 2 mm^2, and their performance, which you typically measure as output jitter, ranges from more than 100 fsec to more than 10 psec.

The wide range of specifications is the result of the wide range of end uses. Some uses include digital-logic or processor clocking, analog-front-end ADC/DAC clocking, serial-link communication, and RF synthesis. This article focuses on frequency-multiplication PLLs, but many other types exist.

PERIOD AND LONG-TERM JITTER

There are many reasons for the difference in power and area among PLLs. The most common reason is the jitter performance, although other requirements, such as output frequency and loop bandwidth, also contribute. Designers should primarily focus on period jitter and long-term jitter. Period jitter is the error that occurs when the output clock itself is acting as the trigger. In this case, you measure jitter at a hold-off time of one output period. In other words, it is the error—that is, phase error—of one clock period. You usually measure period jitter over a large number of samples of the output clock, and you can describe it using a peak-to-peak or an rms (root-mean-square) number.

The period jitter is of concern for static-timing analysis in digital circuits. For example, clocking a digital core at 1 GHz requires a nominal period of 1 nsec. However, no matter how good the PLL is, only the average period is 1 nsec. For static-timing analysis, you must know the shortest period to calculate timing margin. A high-quality PLL has period jitter on the order of 100 fsec for a 1-GHz output. This jitter consumes 0.01% of the output period—orders of magnitude smaller than the uncertainty in static-timing analysis. A PLL with minimal power consumption and area has period jitter on the order of 1 to 10 psec and consumes 0.1 to 1% of the output period, which is usually acceptable.

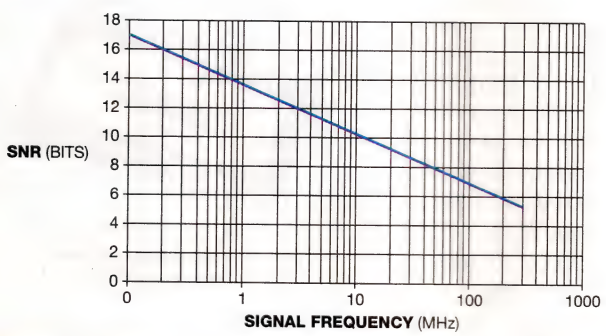

Figure 1 Clocking an ADC with a jittery source worsens the SNR.

Long-term, or N-cycle, jitter is the measure of how much the PLL's output-clock edge deviates from the position of an ideal clock over N cycles, where N is typically thousands of cycles. In other words, long-term jitter is a measure of the accumulated phase error. You usually measure long-term jitter as an rms value rather than a peak-to-peak value.

ALMOST ALL SOCs WITH A CLOCK RATE GREATER THAN 30 MHz USE A PLL FOR FREQUENCY SYNTHESIS.

Long-term jitter is important in applications such as serial-link communications with embedded clocks. These applications include SONET (synchronous optical network), XAUI (10-Gbps attachment-unit interface), and data-converter clocking. For serial-link communications, manufacturers typically specify the long-term jitter at less than 1% rms of a bit period or UI (unit interval). For example, most 10-Gbps serial interfaces specify an rms long-term jitter of less than 1 psec.

For data-converter clocking, the long-term jitter detracts from the SNR (signal-to-noise ratio) because SNR is $1/(2\times\pi\times F\times\sigma)$, where F is the signal frequency, not the sampling frequency, and σ is the rms long-term jitter, which you can assume to be of a gaussian distribution. **Figure 1** provides an example of the SNR versus frequency for an ADC using a clock with 10-psec-rms long-term jitter. High-speed, high-resolution ADCs require precise PLLs. Even 10 psec of rms long-term jitter limits the SNR of an ADC to 10 bits at slightly more than 12 MHz, 12 bits at 3 MHz, and 14 bits at slightly less than 1 MHz.

PLL OPERATION

The operation of the charge-pump PLL involves many

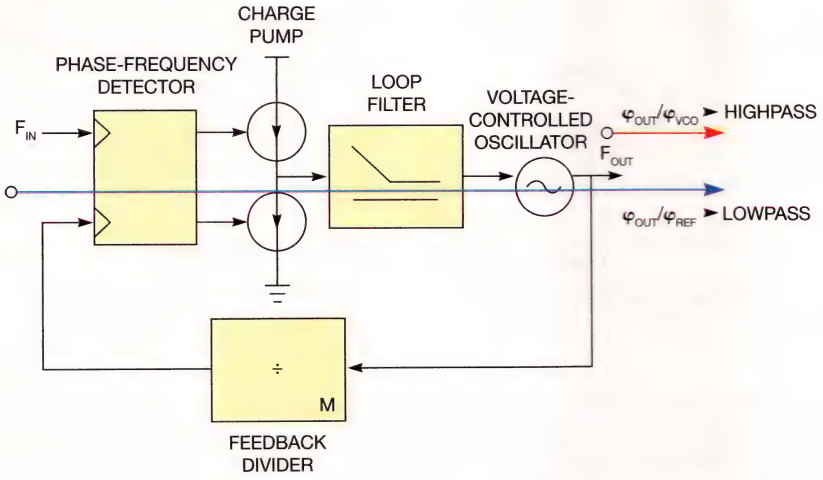

Figure 2 A charge-pump PLL has a lowpass function on the input signal and a highpass function on the phase noise.

trade-offs among jitter, power, and area (**Figure 2**). Many ways exist for implementing PLLs, but most integrated PLLs use this topology. The feedback forces the output frequency, F_{OUT}, to be equal to the input frequency, F_{IN}, multiplied by the feedback-divide value, according to $F_{OUT}=F_{IN}\times M$. Many PLLs also incorporate either an input or an output divider of value N to achieve a frequency $F_{OUT}=F_{IN}\times M/N$.

A detailed frequency-domain analysis concludes that the PLL has both a highpass and a lowpass function (**Reference 1**). There is a lowpass function from input to output, meaning that reference phase noise below the PLL's bandwidth passes through to the output, whereas noise higher than the loop's bandwidth is attenuated. PLLs in noisy environments often use this feature to "clean up" a reference clock by attenuating high-frequency jitter.

The PLL has a highpass characteristic for VCO (voltage-controlled-oscillator) phase noise. Thus, the PLL attenuates VCO phase noise at low frequencies but passes noise above the loop bandwidth to the output. Ideally, all VCO noise would be attenuated through feedback, but PLLs, like all other feedback systems, face bandwidth limitations.

SOURCES OF JITTER

In a typical well-designed PLL, the largest phase noise or jitter source is the VCO. Many other noise sources exist, but you can usually make them smaller than that of a VCO with a modest area or power penalty. The charge pump and loop filter are usually the next-largest noise contributors. The loop filter can be either active or passive. In both cases, most PLLs usually use a zero resistor for loop stabilization. You can also make this noise insignificant by lowering the resistor's value and increasing the integrating capacitor's value and charge pump's current to keep the loop gain constant. This approach has the undesired effect of increasing power and area.

The divider blocks usually contribute negligible device noise. However, a postdivider can be a significant source of short-term jitter due to power-supply noise. Supply noise can also contribute to long-term jitter through the charge pump, loop filter, and VCO, so be sure to design these blocks with sufficient supply-noise rejection.

JITTER AND BANDWIDTH

A frequency-domain analysis shows that jitter is suppressed below the loop bandwidth. The following time-domain experiments show that the PLL bandwidth is the link between short- and long-term jitter. You can perform two time-domain experiments using a signal analyzer or a scope to measure jitter (**Figure 3** and **Reference 2**). The first experiment measures open-loop VCO jitter; the second measures the jitter of a PLL containing the VCO from the first experiment. Both experiments analyze the jitter by measuring the standard deviation of the zero crossings. They measure jitter versus time by using N hold-off times from 1×T to N×T, where T is the nominal period.

The first experiment measures the edges of an open-loop VCO. The standard deviation of the Nth zero crossing is the square root of N times the standard deviation of one cycle ($\sigma_N=\sigma_1\times N^{1/2}$). The standard deviation of one cycle, σ_1, is the period jitter. The value of σ_1 is in practice dif-

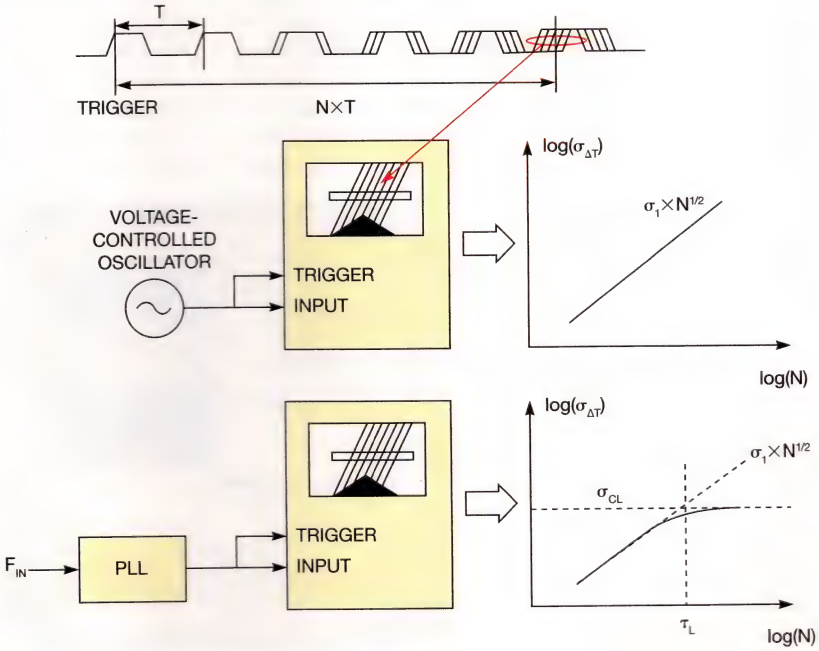

Figure 3 Time-domain experiments measure VCO jitter (top) and jitter from a PLL with the same VCO (bottom).

ficult to measure due to the jitter of any buffer between the VCO and the measurement instrument. The short-term jitter of the instrument itself is also an error source. As N increases, the value of σ_N grows without bound, whereas the rms jitter of the buffers has limits. Therefore, you can extrapolate the value of σ_1 from a plot of σ_N versus N.

A numerical example can highlight the difficulty in directly measuring σ_1. Typical buffer noise for a broadband buffer is on the order of 30 fsec rms. The buffer noise adds in an rms way, so, for example, nine buffers with 30-fsec-rms noise in addition to the VCO with 110-fsec-rms jitter would cause no less than 200-fsec-rms cycle-to-cycle jitter. Additionally, supply noise can be as much as 100 fsec/mV on full-swing buffers, so it may be difficult to measure period jitter of less than 200 fsec in the time domain.

The second experiment measures the edges of a PLL with an ideal reference. The PLL contains the same VCO that the first experiment measured. For a few cycles, the measurements are almost identical to those of the open-loop VCO. You can expect this result because the PLL highpass-filters VCO noise. After many cycles, the measured standard deviation asymptotically approaches the closed-loop standard deviation or long-term jitter, σ_{CL}. The PLL is the force that bounds the phase error.

Figure 3 highlights a few important parameters. The closed-loop parameter σ_{CL} is a function of the PLL's closed-loop bandwidth, τ_L, and the period jitter, σ_1. The open-loop gain, which is a product of the charge-pump, loop-filter, and VCO gain divided by the feedback-divide value, determines the closed-loop bandwidth, a system-design parameter. You can calculate the closed-loop bandwidth, normalized to one VCO period, T, as $1/(2\pi F_L/F_{VCO})$. You can now calculate the long-term jitter as $\sigma_{CL}=\sigma_1/(4\pi F_L/F_{VCO})^{1/2}$ (**Reference 2**).

This analysis is a simplification in at least two ways. First, the only noise source it considers is VCO phase noise. However, VCO noise limits most well-designed PLLs. Note that this analysis does not consider supply noise or reference noise. The second simplification is that this analysis assumes the PLL to be a first-order loop. Most PLLs are at least second-order loops. Many PLLs are overdamped, however, and appear almost as first-order loops for the sake of this analysis. Additionally, the long-term jitter is a function of the square root of the bandwidth, so the error is not too severe for the sake of first-pass manual calculations.

These experiments yield two important results. The first result is that short-term period jitter depends almost entirely on the VCO and output buffers and does not depend on the PLL bandwidth. The second result is that long-term jitter depends on both the VCO's and the PLL's bandwidth, and it improves as the VCO improves and the bandwidth increases.

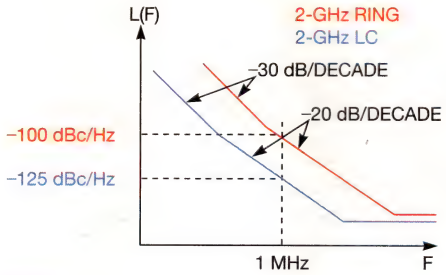

Figure 4 VCO phase noise has three distinct regions.

VCO PHASE NOISE

A pair of equal-power, 2-GHz oscillators consume several milliwatts (**Figure 4**). One oscillator is ring-based, and the other is LC-based. Three distinct regions of operation are shown in **Figure 4**. The most important is the –20-dB/decade region. This region typically determines the period jitter of the VCO, σ_1.

The plot also labels the –30-dB/decade region of the VCO. In this region, flicker device noise typically dominates over white noise, causing the increased slope. Because flicker noise is responsible for the increased slope, the transition from –30 dB/decade to –20 dB/decade is the flicker corner of the VCO. For ring-based VCOs, the flicker-noise corner typically ranges from 300 kHz to 3 MHz. For LC-based VCOs, you can obtain flicker-noise corners of less than 100 kHz. You should take care to optimize the VCO for flicker noise (**Reference 3**).

The plots also include a flat region at high frequencies, due to VCO output buffers. This region is important for period jitter but typically not important for long-term jitter, as the following **equation** shows: $L_{dB}(F)\approx 10\times\log_{10}[(1/P_{SIG})\times(F_{OSC})^2/(Q\times F)^2]$. From this **equation**, the phase noise drops by 3 dB for a twofold increase in power for a given oscillator frequency. Increasing the power can be an effective way of improving the phase-noise performance but can become expensive. A 20-dB improvement in phase noise comes at a cost of a 100-fold power increase, with all other things constant. Another way of improving the phase noise is to increase the quality factor of the resonant tank. Doubling the quality factor halves the phase noise, a 6-dB improvement. The inductor structure often limits the achievable quality factor in a CMOS process. Typical quality factors range from seven to 15 and vary with many factors, including frequency and metal thickness. A trade-off also exists between tuning range and quality factor in LC VCOs in which a higher quality comes at the expense of a smaller tuning range.

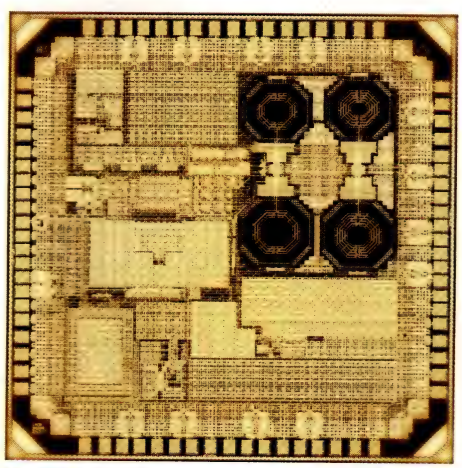
Figure 5 This PLL die uses an LC-based oscillator and provides excellent jitter specs.

There is roughly a 20-dB difference between the phase noise of a typical ring oscillator and that of an LC oscillator of the same power in a deep-submicron CMOS processes. This difference illustrates the phase-noise advantage of resonant-tank structures for phase noise.

As noted, the value of σ_1 can be difficult to measure in the time domain. However, it is relatively straightforward in the frequency do-

main. You can calculate the VCO's period jitter, σ_1, as $\sigma_1^2=F^2\times L(F)/F_{OSC}^3$, where F is the offset frequency, L(F) is the phase noise at F, and F_{OSC} is the frequency of oscillation (**Reference 4**). In this example, the period jitter for the ring VCO with −100 dBc/Hz at 1-MHz offset and 2-GHz oscillation frequency is 112 fsec rms. The LC oscillator with −125 dBc/Hz at 1-MHz offset and 2-GHz oscillation frequency results in σ_1 of 6.3 fsec rms. These values are typically too small to directly measure in the time domain, and buffer and scope noise obscure them.

You can calculate the long-term jitter from the PLL bandwidth and the corresponding value of σ_1 according to $\sigma_{CO}=\sigma_1/(4\pi F_L/F_{VCO})^{1/2}$. Again, the calculations assume an overdamped PLL, VCO noise only, and no supply noise. Assuming a bandwidth of 100 kHz, the ring PLL with a σ_1 of 112 fsec would have approximately 4.5 psec rms of long-term jitter, whereas the LC PLL with a σ_1 of 6.3 fsec would have a long-term jitter of 270 fsec rms. If you increase the bandwidth to 1 MHz, then both long-term jitter values would decrease by $\sqrt{10}$ to 1.4 and 85 fsec, respectively. You could continue these calculations for higher and higher bandwidths, but many reasons exist for limiting the bandwidth, and the jitter would not continue to decrease.

One of the primary factors limiting bandwidth is the PLL's stability. Bandwidth is typically approximately only 1/20th of the reference rate for adequate phase margin. For high-performance PLLs, a low loop bandwidth mitigates reference-clock feedthrough. Suppressing reference-clock spurs typically requires a bandwidth of no more than 1/100th of the reference rate. Other reasons to limit the PLL bandwidth include delta-sigma modulation and reference-noise, loop-filter, and charge-pump noise suppression.

PLL AREA

Along with performance and power, area is a major specification for PLLs. The performance level of a PLL largely determines its area. You gain a large increase in performance by choosing an LC-based VCO rather than a ring-based VCO. The LC-based VCOs usually measure at least 300×300 microns for a design with one inductor and can be even larger. Ring oscillators, on the other hand, can measure 40×40 microns or smaller. Typically, LC oscillators have a narrower tuning range than do ring oscillators. Therefore, you must sometimes use multiple VCOs in the same PLL to achieve a wide tuning range, further increasing the area.

The loop filter is another part of the PLL that significantly increases in area with performance. An integrated loop filter can consume 500×500 microns or more. As the performance of the PLL drops, you can scale down the resistors, capacitors, and charge-pump current to reduce area at the cost of noise. You can make a SONET/multiprotocol clocking IC in 0.13-micron CMOS (**Figure 5**). The **figure** clearly shows the four-core LC VCO. The PLL measures roughly 1.4 mm^2 in area. The long-term jitter is less than 500 fsec rms with a 50-kHz bandwidth. The PLL's power dissipation is approximately 70 mW, depending on the mode of operation.

Figure 6 This PLL die uses a ring-based oscillator and is one-tenth the size and power of an LC-based oscillator—at the expense of worse jitter performance.

You can make a ring-based fractional-N PLL in 0.13-micron CMOS (**Figure 6**). The PLL's area is 0.09 mm^2, which is more than 10 times smaller than the LC PLL example in **Figure 5**. The long-term jitter is as low as 3 psec rms with a 1-MHz bandwidth, and the power consumption is approximately 5 mW, depending on the mode of operation. The area is largely digital. The digital blocks include a delta-sigma modulator, a predivider, a postdivider, feedback dividers, and control circuits. The analog area is more than 10 times smaller than that of the LC PLL analog area.

The two PLLs in **figures 5** and **6** show why a one-size-fits-all approach does not work for SOC PLLs. The first PLL has jitter that is sufficient for almost all SOC applications. However, the area and power are both more than 10 times the area and power of the second example. The long-term jitter of the second example is six times higher, however, and would be almost 20 times higher if the PLLs had the same bandwidth.

The driving factor for many of the trade-offs of PLL SOCs is long-term jitter. If the long-term-jitter specification is loose, you can use a small, low-power, ring-based PLL. A tighter long-term jitter specification entails the use of a lot of silicon area and power to meet the requirement with an LC-based PLL. However, for many applications between the two extremes, the choice is not clear, and you must perform careful analysis to optimize the PLL for power and area. EDN

REFERENCES

1. Garner, Floyd M, *Phaselock Techniques*, Third Edition, John Wiley & Sons, 2005.
2. McNeill, JA, "Jitter in ring oscillators," *IEEE Journal of Solid-State Circuits*, Volume 32, June 1997, pg 870.
3. Hajimiri A, and TH Lee, "A general theory of phase noise in electrical oscillators," *IEEE Journal of Solid-State Circuits*, Volume 33, February 1998, pg 179.
4. Poore, R, "Phase Noise and Jitter," Agilent Technologies, May 2001.

AUTHORS' BIOGRAPHIES

Jeff Galloway is a co-founder of Silicon Creations and is responsible for analog IP (intellectual-property) design and development. Before founding Silicon Creations, he held various positions at Hewlett-Packard, Agilent Technologies, and Mosaid. Galloway has a bachelor's degree in electrical engineering from the Georgia Institute of Technology (Atlanta) and a master's degree in electrical engineering from Stanford University (Palo Alto, CA).

Randy Caplan is co-founder and vice president of engineering at Silicon Creations. He is responsible for PLL-product architecture and design, as well as the technology road map for the company. Before founding Silicon Creations, he held various positions at Agilent Technologies, Virtual Silicon, and Mosaid. Caplan has a bachelor's degree in electrical engineering from the Georgia Institute of Technology (Atlanta).

Tiny Octal DAC with 10ppm/°C Reference

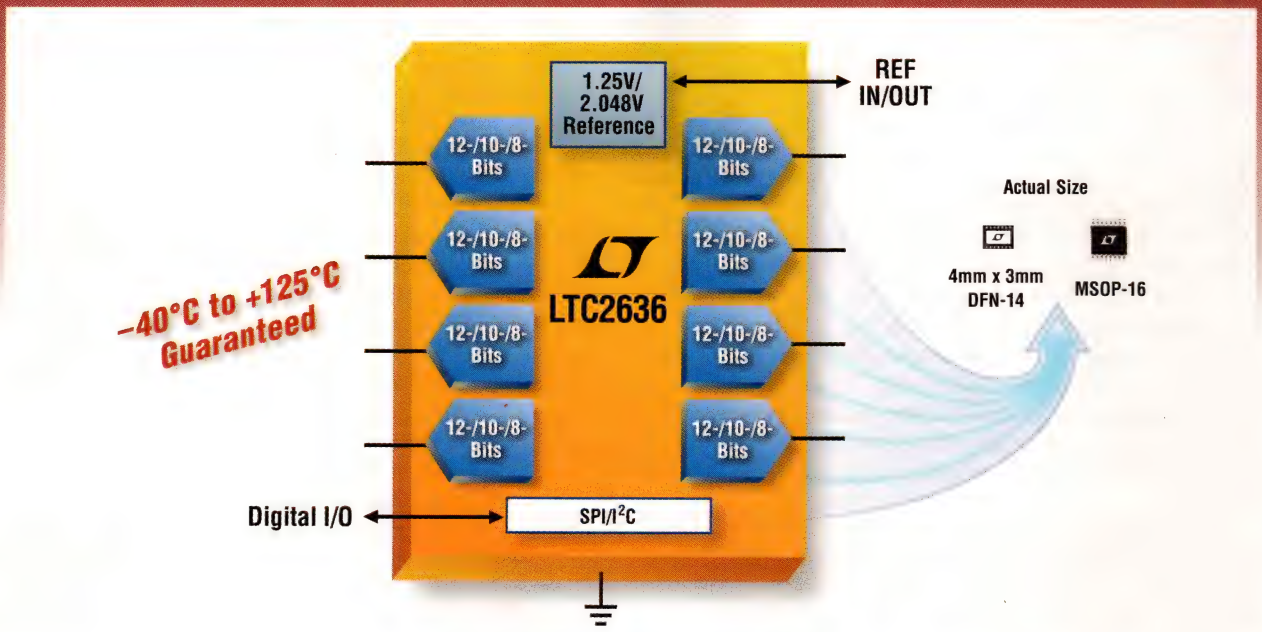

12-, 10- & 8-Bit DACs Reduce Solution Size by Half

Our LTC®2636 DAC family includes singles, duals, quads and octals with ordering options for 12-, 10- or 8-bits of resolution, internal 1.25V or 2.048V 10ppm/°C reference, SPI or I²C interface, power-on reset to zero- or mid-scale, and commercial, industrial or automotive guaranteed temperature specifications. The smallest solution available, these DACs are ideal for industrial, optical networking and automotive applications. Customize your design today.

Features

- Internal Precision Reference: 1.25V or 2.048V, 10ppm/°C
- 2.7V to 5.5V Supply Range
- Ultralow Crosstalk Between DACs (<2.4nV•s)
- –40°C to +125°C Guaranteed Temperature Range
- Power-on Reset to Zero-Scale or Mid-Scale
- 16-Pin MSOP, 4mm x 3mm DFN-14 Packages

Multichannel DAC Family

Part Number	Bits	I/O	DACs
LTC2636	12/10/8	SPI	8
LTC2637	12/10/8	I²C	8
LTC2634	12/10/8	SPI	4
LTC2635*	12/10/8	I²C	4
LTC2632*	12/10/8	SPI	2
LTC2633*	12/10/8	I²C	2
LTC2630/ LTC2640	12/10/8	SPI	1
LTC2631	12/10/8	I²C	1

*Future Products. Contact factory for availability.

Info & Free Samples

www.linear.com/2636

1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

LEDs light LAN-cable tester

Noureddine Benabadji, University of Sciences and Technology, Oran, Algeria

Cable and connector faults are relatively common in LANs (local-area networks). You can use the circuit in **Figure 1** to test straight- or cross-link 10BaseT, 100BaseT, or gigabit UTP (unshielded-twisted-pair) and STP (shielded-twisted-pair) cables. The circuit performs a continuity test for each linked pair. An LED representing each linked pair flashes when the corresponding pair connects properly between the cable's RJ-45 connectors. It can also determine whether the cable is a straight-link or a cross-link type.

IC_1, a small, six-pin Microchip (www.microchip.com) PIC10F200 microcontroller, performs the test. On power-up, the four I/O pins act as outputs and are driven high for approximately 0.5 second and then are driven low in the following sequence: GP0, GP1, GP2, and GP3. Once this task finishes, the microcontroller has a dead time for about 4 seconds, after which it loops to the beginning to repeat the procedure.

Adding a low-power BS170 N-channel MOSFET allows Pin 6 (GP3), normally an input-only pin, to act as an output, according to a previous Design Idea (**Reference 1**). The power-supply voltage of 5V is not critical, and you can choose any supply from 2 to 5.5V for the 8-bit PIC microcontroller. **Figure 2** shows cable wiring for both types of ca-

DIs Inside

See all of *EDN*'s Design Ideas at www.edn.com/designideas.

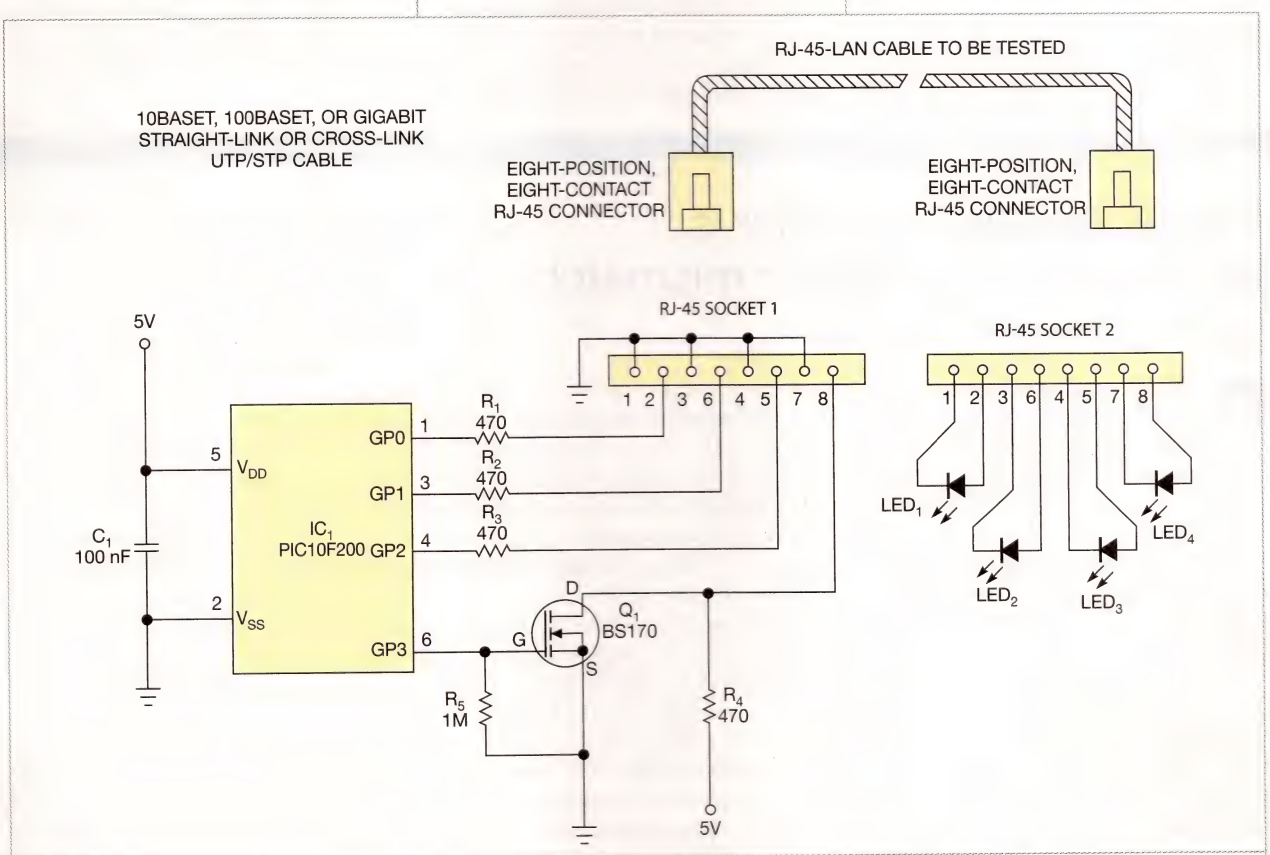

Figure 1 You can use this circuit to test straight- or cross-link 10BaseT, 100BaseT, or gigabit UTP and STP cables.

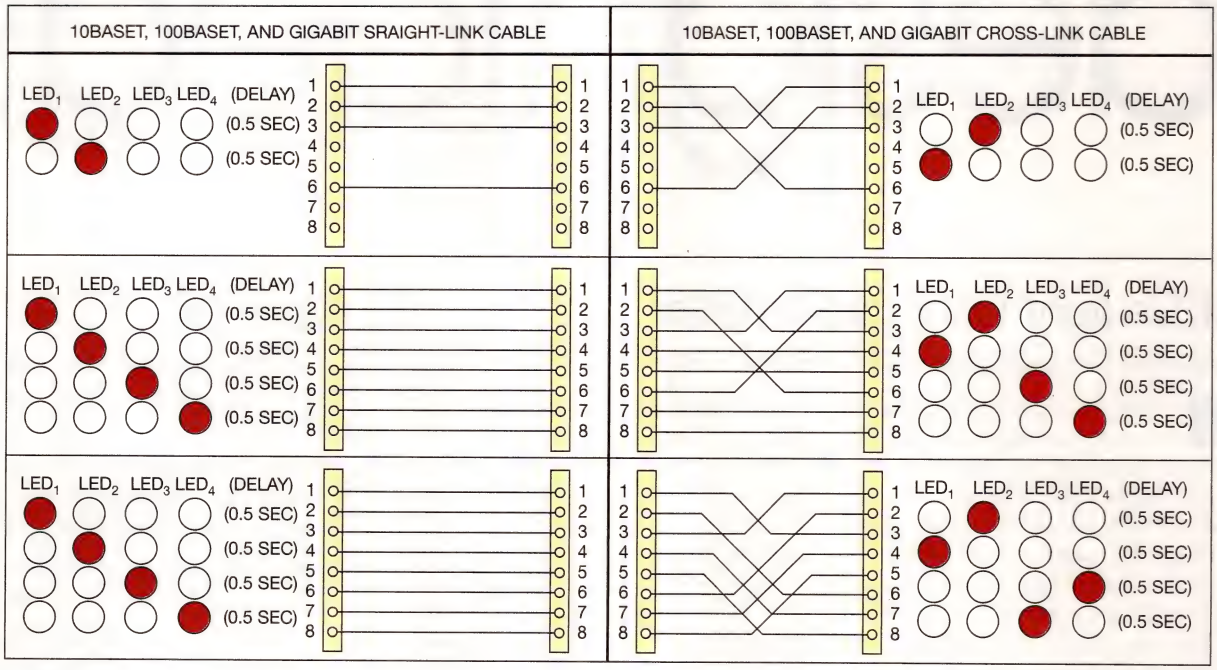

Figure 2 Cable wiring for both straight-link and cross-link cable fits with RJ-45 wiring.

bles and how they fit with RJ-45 connectors. For each type, the microcontroller initiates lighting. You can mount the RJ-45 socket in a separate small box to test the inside wall's long cables. You can view the complete and fully commented source code in the Web version of this Design Idea at www.edn.com/101202dia. It uses fewer than 256 words of memory. Because it does not specifically target the PIC10F200, it is easy to understand and adapt to other microcontrollers. EDN

REFERENCE

1 Muñoz, Antonio, "Use the MCLR pin as an output with PIC microcontrollers," *EDN*, Jan 10, 2008, pg 65, www.edn.com/article/470769-Use_the_MCLR_pin_as_an_output_with_PIC_microcontrollers.php.

Instrumentation amplifier is less sensitive to resistor mismatch

Reza Moghimi, Analog Devices, San Jose, CA

This Design Idea offers a new instrumentation-amplifier topology that improves ac CMR (common-mode rejection). Because it uses discrete amplifiers, you can customize it for the lowest power, price, noise, and supply voltage, depending on your application's requirements. Previously, instrumentation amps using discrete components had poor CMR. System designers still build their own using discrete designs because standard integrated instrumentation amps don't meet their design requirements, have undesired package options, or are too expensive. You can build discrete instrumentation amps with two or three op amps and a few resistors (**Figure 1**). Most monolithic instrumentation amps, however, use the three-op-amp configuration. This approach provides the best ac and dc CMR. A big challenge when building discrete instrumentation amps is to achieve CMR on par with that of monolithic instrumentation amps.

The poor CMR of the three-op-amp topology (**Figure 1b**) is due to resistor mismatching. Any common-mode signal at V_{IN1} and V_{IN2} appears at the outputs of A_1 and A_2. The difference amplifier, comprising R_1, R_2, R_3, R_4, and A_3, should reject this common-mode voltage. To achieve high CMR, this difference amplifier requires matched resistors and the A_3 op amp, which specifies high CMR. With 0.1%-resistor matching, the best possible CMR at dc is 54 dB. This CMR further degrades over frequency, depending on the op amp you select.

The poor CMR in two-op-amp instrumentation amps is due to the unequal phase shift that occurs at the two inputs of A_2. Signal V_{IN1} must travel through amplifier A_1 before amplifier A_2 subtracts it from V_{IN2}. Thus, the voltage at the output of A_1 becomes delayed with respect to V_{IN1}. Amplifier A_1 introduces the delay that causes its output to lag behind the directly applied input voltage at V_{IN2}. This phase difference results in an instantaneous difference in A_1's output and V_{IN2}, even if the amplitudes of both voltages are equal. This difference causes a frequency-dependent common-mode error voltage at V_{OUT}. Further, this ac common-mode error increases with common-mode frequency because the phase shift through A_1 increases with frequency due to the single-pole roll-off frequency response.

Instrumentation-amplifier-design guidelines give the error relationship,

Need for speed?

Industry's smallest 2.5A step-down converter with I²C control

The MAX8952 is a tiny, DC-DC step-down converter with high output current that allows application processors to run at their maximum clock speed to meet the demands of today's multimedia applications in handheld devices. Designed to minimize your solution area, this 2mm x 2mm converter also ensures tight output-voltage accuracy for reliable, power-efficient operation.

Feature	Competitor 1	Competitor 2	MAX8952	MAX8952 advantage
I²C control	No	No	Yes	Enables dynamic voltage scaling without external components
Frequency (MHz)	3	1.2	3.2	Smaller solution size
Package size (mm x mm)	2 x 2	4 x 4	2 x 2	
Solution size (mm²)	43.28	122.8	40.5	
Initial accuracy* (%)	±1.5 (±18mV)	±1 (±12mV)	< 0.5	Enables reliable operation and lower power consumption
Load/line/temp** (%)	≥ ±1.5	≥ ±3	±0.71	
I_{OUT} (A)	2	2	2.5	Run your processor at maximum speed

*Conditions: V_{IN} = 3.6V, V_{OUT} = 1.2V
**Conditions: V_{IN} = 3.6V to 5V, V_{OUT} = 1.2V, I_{OUT} = 0.5A to 2A, temperature range = -40°C to +85°C

www.maxim-ic.com/MAX8952-info

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM
INNOVATION DELIVERED®

For free samples or technical support, visit our website.

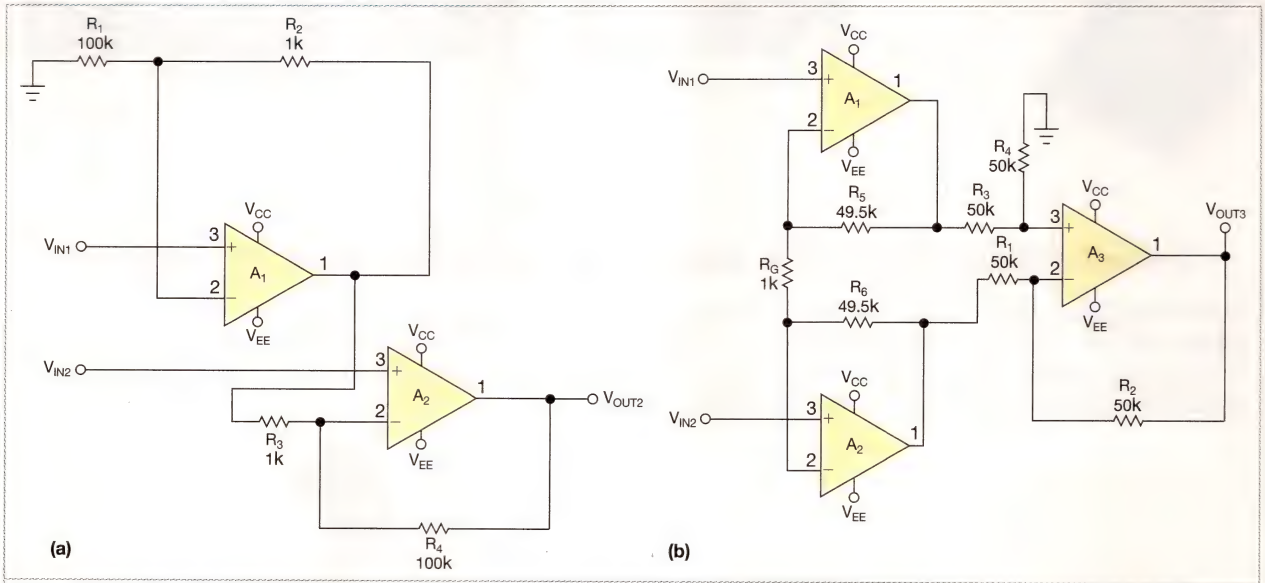

Figure 1 You can use conventional two-op-amp (a) and three-op-amp (b) configurations to build an instrumentation amplifier.

according to the following **equation**: (100 Hz/100 kHz)×100%=0.1%, when you design an instrument amplifier for a closed-loop bandwidth of 100 kHz by using micropower op amps with a unity-gain bandwidth of 500 kHz configured for a gain of five (**Reference 1**). A common-mode error of 0.1% is equivalent to a 60-dB CMRR (common-mode-rejection ratio). Even if you trim this circuit to achieve 100-dB CMRR at dc, this performance would hold only for frequencies less than 1 Hz. At 100 Hz, the CMRR could not be better than 60 dB.

You can overcome the shortcomings of the two-op-amp circuit if you compensate for the phase delay of A_1 using active-feedback compensation. You can reduce phase error to virtually zero at low frequencies with a dual matched op amp (**Reference 2**). This approach adds an equal but opposite phase shift in the amplifier's feedback loop. Because the circuit depends on amplifier matching, you must use a dual- or quad-op-amp part.

The active feedback requires an extra op amp and two external resistors to achieve phase-error cancellation (**Figure 2**). In the circuit, amplifier A_1 provides forward gain of the composite amplifier. K_1 determines the closed-loop gain (A_V=1+K_1), where A_V is voltage gain and K is a constant. Amplifier A_2 provides feedback to amplifier A_1. K_2 determines the amount of phase-error compensation and has no effect on the forward gain of the composite amplifier. Optimum error cancellation occurs when K_1 is equal to K_2. The error terms are functions of complex frequency response, according to the following **equations**:

$$A_{ERROR}=\left(\frac{\omega_{CLOSED\_LOOP}}{\beta\times\omega_{VGB}}\right)^2;$$

$$\phi_{ERROR}=\left(\frac{\omega_{CLOSED\_LOOP}}{\beta\times\omega_{VGB}}\right)^3.$$

Figure 3 shows the new configuration, which applies an active feedback network to a two-op-amp configuration.

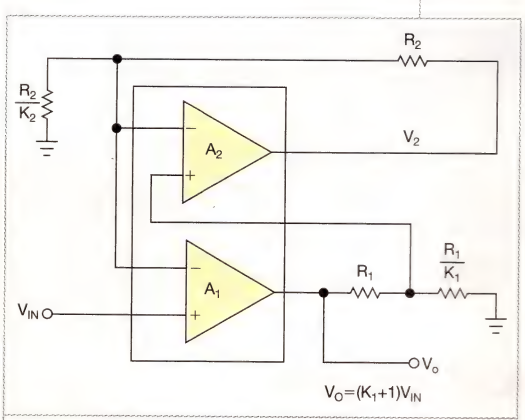

Figure 2 Using A_2 in an active-feedback network greatly improves A_1's effective bandwidth. The op amps must match, so you should use a dual-amplifier package for this circuit.

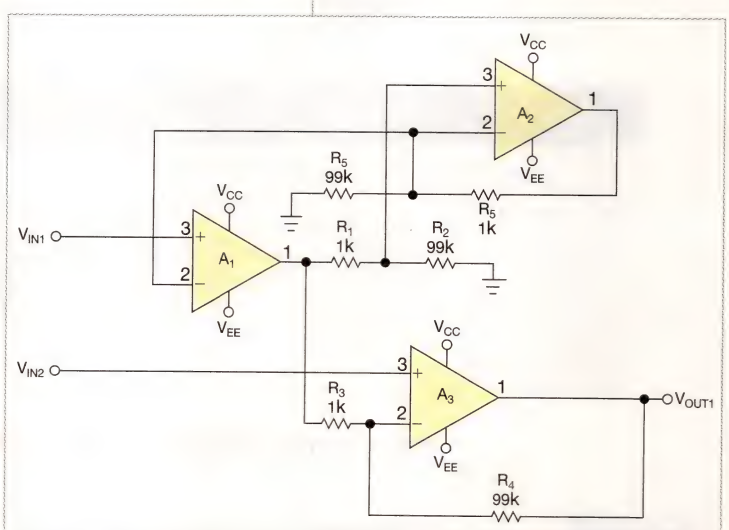

Figure 3 The three-op-amp instrumentation amplifier employs a two-amplifier topology; a third amplifier provides active-feedback phase compensation.

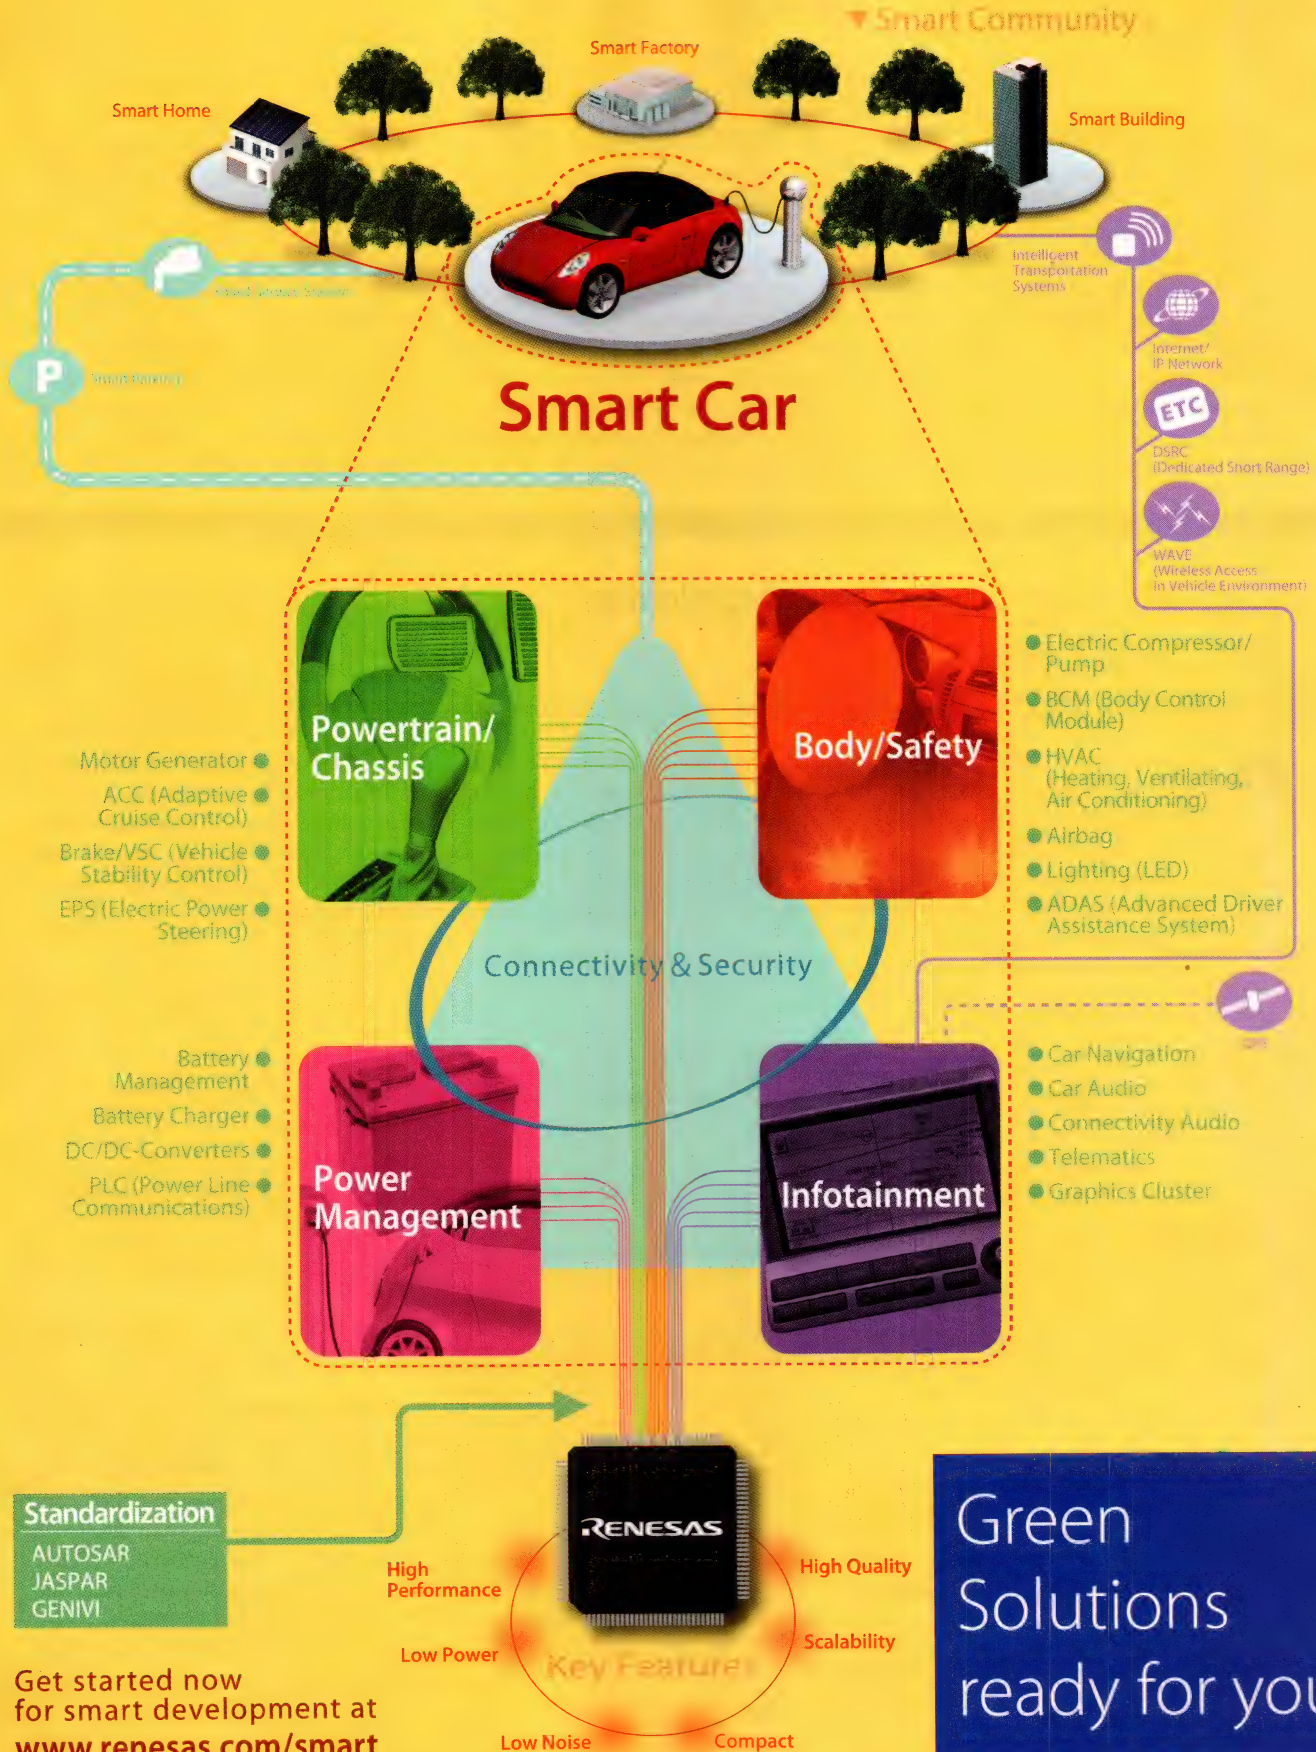
Smart Community
Smart Factory
Smart Home
Smart Building
Intelligent Transportation Systems
Internet/ IP Network
ETC
DSRC (Dedicated Short Range)
WAVE (Wireless Access in Vehicle Environment)
Smart Car
Powertrain/ Chassis
Body/Safety
Motor Generator
ACC (Adaptive Cruise Control)
Brake/VSC (Vehicle Stability Control)
EPS (Electric Power Steering)
Electric Compressor/ Pump
BCM (Body Control Module)
HVAC (Heating, Ventilating, Air Conditioning)
Airbag
Lighting (LED)
ADAS (Advanced Driver Assistance System)
Connectivity & Security
Battery Management
Battery Charger
DC/DC-Converters
PLC (Power Line Communications)
Power Management
Infotainment
Car Navigation
Car Audio
Connectivity Audio
Telematics
Graphics Cluster
Standardization
AUTOSAR
JASPAR
GENIVI
RENESAS
High Performance
High Quality
Low Power
Scalability
Key Feature
Low Noise
Compact
Green Solutions ready for you
Get started now for smart development at
www.renesas.com/smart
Renesas Electronics Corporation
RENESAS

You correct the phase delay through A_1 using the active-feedback circuitry comprising A_2 and four resistors. The classic two- and three-amplifier instrumentation amplifiers and this proposed three-amplifier topology using active feedback have similar dc CMRRs. However, the circuits' ac CMRs differ. This difference in performance arises when you apply and sweep a large ac common-mode voltage at the input. These test specs are for a circuit using Analog Devices' (www.analog.com) AD8603 op amps. The supply voltage is ±2.5V, the common-mode voltage is 0.001 mV, and the gain is 100. You take error measurements with a 500-Hz input signal.

You then change the operating conditions to a common-mode voltage of 2.001V and a gain of 100. The circuit in **Figure 3** provides great improvement over the circuit in **Figure 1a**. The circuit has less sensitivity to resistor mismatch than the circuit in **Figure 1b** with perfectly matched resistors. (For more data, go to www.edn.com/101202dib.)EDN

REFERENCES

1 Kitchin, Charles, and Lew Counts, *A Designer's Guide to Instrumentation Amplifiers*, Chapter 2, Analog Devices, 2000, pg 2-1, http://bit.ly/b0rXnT.

2 Wong, James, "AN-107 Active Feedback Improves Amplifier Phase Accuracy," Analog Devices, September 1987, http://bit.ly/a8OI38.

Circuit keeps wandering children and pets nearby

Tom Au-Yeung and Craig Sakamoto, Maxim Integrated Products, Sunnyvale, CA

The receiver circuit in **Figure 1** sounds an audio alarm when the transmitter (**Figure 2**) moves beyond a designated perimeter. The transmitter, a voltage-controlled oscillator, operates at approximately 915 MHz in the unlicensed ISM (industrial/scientific/medical) band. It has a tuning voltage of $1.5V=3\times R_2/(R_1+R_2)$, which lets you easily adjust the frequency by varying the values of resistors R_1 and R_2.

The receiver comprises low-noise amplifier IC_1, power detector IC_2, comparator IC_3, and a buzzer. When the transmitter is within range—for example, when a child or a pet is carrying it—the receiver detects the RF signal and provides a voltage greater than 400 mV at the inverting terminal of the comparator. Resistors R_9 and R_{10} preset the reference voltage at the comparator's noninverting terminal. The reference voltage is $3\times R_{10}/(R_9+R_{10})$, and the comparator's output remains low.

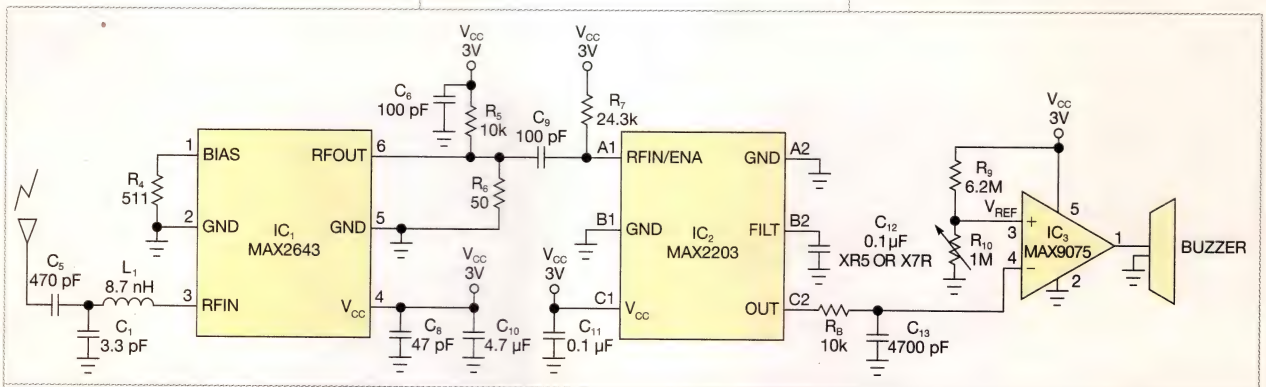

Figure 1 This 915-MHz receiver sounds an alarm when the comparator's inverting-input voltage drops below 400 mV.

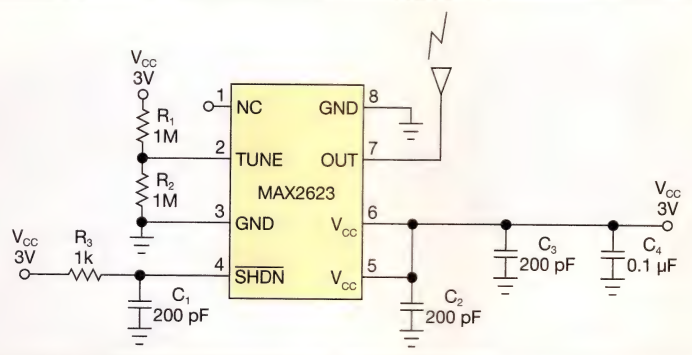

Figure 2 The transmitter comprises a voltage-controlled oscillator, which R_1 and R_2 tune to approximately 915 MHz.

When the transmitter moves outside the predetermined boundary, the detected RF produces less than 400 mV at the comparator. The comparator then generates an output of approximately 3V, which turns on the buzzer and sounds an alert that the transmitter has moved beyond the restricted perimeter. To increase the detection range, you can place additional low-noise amplifiers or VGAs (variable-gain amplifiers) in front of the power detector. You can also increase or decrease the desired perimeter by adjusting R_{10} to change the comparator's reference voltage.EDN

The West Coast's Project Acceleration Resource for Electronics Manufacturing

Electronics® WEST

February 8–10, 2011

Anaheim Convention Center
Anaheim, CA

Your face-to-face meeting place for the latest technologies in electronics manufacturing, including:

- Assembly Board Inspection
- Components
- Contract Services
- Production Machinery
- Software
- Subassemblies
- Test & Inspection Products

. . . and much more

Your Project Acceleration RESOURCE

Supported by:

EDN

CIRCUITS ASSEMBLY
The Journal for Surface Mount and Electronics Assembly®

Submit your registration online today at
ElectronicsWestShow.com
Please use Promo Code AA

Produced and Managed by Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA, 90064-1549 Tel: 310/445-4200 cancom.com

"I need a faster, more accurate audio analyzer."

We knew you'd say that.

The new Agilent U8903A audio analyzer is four times more accurate and over seven times faster than its popular HP 8903B ancestor. In fact, it is among the fastest and most accurate currently in its class. It's also a more fully capable audio analyzer, letting you make crosstalk measurements and phase measurements, FFT analysis, and graph sweep. That's taking it to the next level. That's Agilent.

Learn how to migrate effortlessly - Free application note filled with tips
www.agilent.com/find/AgilentU8903A

u.s. 1-800-829-4444 canada: 1-877-894-4414

Agilent Technologies

Electronics industry braces for rare-earth-materials shortages

China has started to severely restrict the exports of rare-earth materials, which often find use in "green"-technology designs, including hybrid vehicles and energy-efficient lighting, as well as in the medical, defense, and consumer markets. The country delivers nearly 100% of the world's rare-earth materials: 17 metals that are somewhat hard to refine and that tend to occur in the same ore deposits (**Table 1**). The cutbacks have resulted in shock waves through the electronics industry and could force design changes in the near future.

China set out on a moderate restriction path this year and then announced in July that it would cut exports by 72% for the remainder of 2010. It plans an overall export reduction of 30% for next year.

These cutbacks have increased the price of rare-earth materials an average of 700%, prompting legislation, which is currently stalled, to restart US rare-earth-materials production. The Western Hemisphere's one rare-earth-materials producer, Colorado-based Molycorp Minerals (www.molycorp.com), issued an initial public offering of stock in July, raising $390 million to restart its California mine and ramp up processing to counter world shortages.

Production of rare-earth materials fell off worldwide beginning in the 1980s when low prices in China made production unfeasible elsewhere in the world. Tom Valiere, senior vice president and co-founder of Design Chain Associates (www.designchainassociates.com), says this cutback is a wake-up call for US industry. "We used to lead the world in the export of rare-earth materials," he says. "In the last 20 years, we've become dependent. The whole thing flew under the radar until green technology placed demand on rare-earth materials and we realized they were sole-sourced to China."

China's restrictions this year have been part of a multiyear plan to save most of its supply for its own industry. "Each year, China has brought down its exports by X% and never exceeded its quotas," says Gareth Hatch, co-founder of Technology Metals Research (www.techmetalsresearch.com). "The reduction the country made in July was a huge reduction over the first half of the year."

Worldwide shortages are now occurring. "The world outside China uses a collective 50,000 tons annually," says Jim Simms, director of public affairs at Molycorp Minerals. "[China] reduced its exporting in 2010 to about 30,000 tons. Since China supplies about 97% of rare-earth materials, the world has to depend on what China exports."

Simms believes that the demand for the materials will just increase over the coming years. The company expects to produce 20,000 tons by the end of 2012. "My BlackBerry only has about 3/10g of rare-earth materials," he says, but "a single wind turbine requires about one ton. A car can use about 25 kg."

Rare-earth materials include terbium, which finds use in flat-panel TVs and high-efficiency fluorescent lamps, and neodymium, key to the permanent magnets in high-efficiency electric motors. Rare-earth materials are not indeed rare. The series of nonferrous metals is common in the environment. According to Design Chain Associates, most rare-earth materials are as common as copper, and even the rarest is more common than gold.

A US Government Accountability Office report claims that China's dominance in rare-earth materials' mining and refining has implications for global availability and pricing and could also jeopardize US defense readiness. Meanwhile, China's export restriction of rare-earth materials has hit Japan the hardest, especially for its production of high-efficiency electric motors. Nidec Corp, a maker of electric cars, has developed a process, starting in 2012, to build switched-reluctance motors without rare-earth materials.

—by Rob Spiegel

TABLE 1 RARE-EARTH MATERIALS

Periodic no.	Symbol	Name
21	Sc	Scandium
39	Y	Yttrium
57	La	Lanthanum
58	Ce	Cerium
59	Pr	Praseodymium
60	Nd	Neodymium
61	Pm	Promethium
62	Sm	Samarium
63	Eu	Europium
64	Gd	Gadolinium
65	Tb	Terbium
66	Dy	Dysprosium
67	Ho	Holmium
68	Er	Erbium
69	Tm	Thulium
70	Yb	Ytterbium
71	Lu	Lutetium

productroundup

CONNECTORS

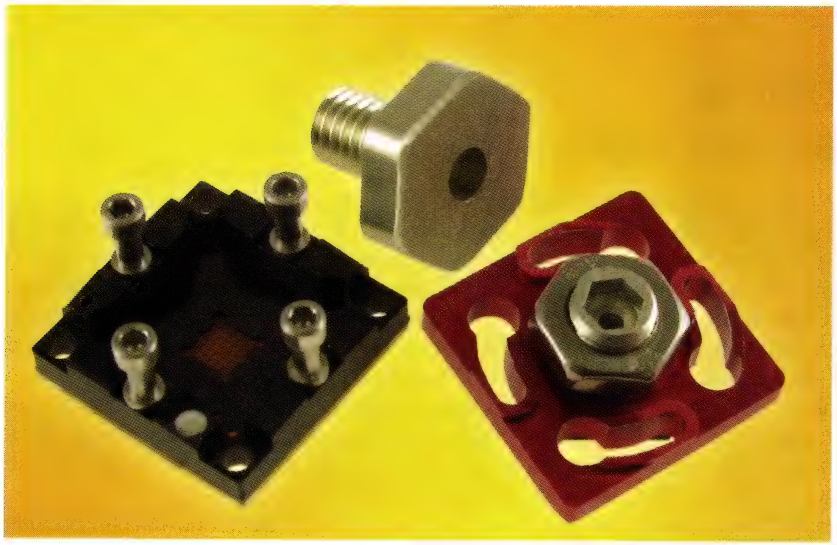

WCSP socket has small footprint

The SG-BGA-7188 WCSP (wafer-chip-scale-package) socket targets use in 0.5-mm-pitch, 25-pin WCSPs. The socket fits into 2.75×2.75-mm packages and operates at bandwidths as high as 10 GHz with less than 1 dB of insertion loss. The contact resistance is typically 20 mΩ per I/O. The device connects all pins with 10-GHz bandwidth on all connections. The socket is mounted using supplied hardware on the target PCB with no soldering. The socket also incorporates a quick-insertion method using an integrated compression plate to the socket lid so that users can quickly change out ICs. The socket verifies the function of ICs in a development system that is the power-management unit in an embedded camera module. The device also features a replaceable IC guide that allows testing of other package sizes with 0.5-mm pitch, such as 2.25×2.25, 2.25×2.75, and 1.75×2.25 mm. The pin's self-inductance is 0.15 nH, and mutual inductance is 0.025 nH. Capacitance to ground is 0.01 pF, and current capacity is 2A/pin. The SG-BGA-7188 sells for $221.34 (100).

Ironwood Electronics, www.ironwoodelectronics.com

Battery retainer has narrow profile

The BK-913 CR2032 lithium battery retainer targets use in 20-mm-diameter batteries. The device is 23% narrower than competing designs and weighs 0.8g, making it suitable for greeting cards, consumer products, electronic-toll tags, single-use devices, RFID, and other price-sensitive applications. The nickel-plated phosphor-bronze retainer has two PC pins. It features dual pressure contacts that offer low electrical resistance and keep a tight connection with the battery. Shock and vibration specs comply with EIA-364, and users can easily change the retainer without tools. The BK-913 uses BR2032 and CR2032 lithium coin cells. Rechargeable types include ML2032, VL2032, and LIR2032 batteries. The BK-913 also accepts other batteries with a height of 2.5 to 3.2 mm. Other retainers in the family accept 4- to 24-mm-diameter cells. The retainers sell for 12 cents (10,000).

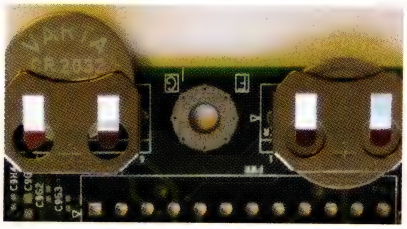

Memory Protection Devices, www.memoryprotectiondevices.com

SuperSpeed USB 3.0 receptacles have through-hole terminations

The USB3-A and -B series of SuperSpeed USB 3.0-certified A- and B-type receptacles, respectively, have through-hole terminations. The mating USBC3 series is available with an A-type plug on one end and a B-type plug on the other end in a choice of wire lengths. The UUSB3 series micro SuperSpeed USB 3.0 version is also available. Prices start at approximately 44 cents per connector (5000).

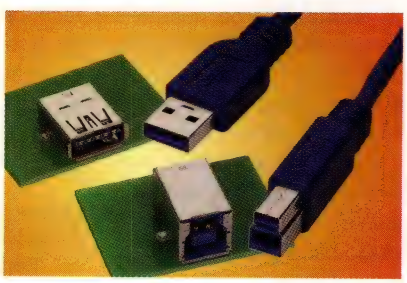

Samtec, www.samtec.com

Rugged connector is water-resistant

This series of water-resistant D-sub connectors has rugged, one-piece solid-body design, performing better against mechanical impact during installation and operation. With fewer leakage paths, the connectors target use in industrial machinery, medical equipment, test-and-measurement equipment, and communications. The

IP67-rated connectors have a silicone O-ring panel seal for watertight connection as well as gold-plated machined

contacts. They come in standard, high-density, and combination D-sub configurations and with solder-cup, right-angle, and straight terminations. The connectors ensure a continuously secure connection and are dustproof. Prices start at $6 (OEM quantities).

Conec, www.conec.com

Automotive socket is weather-resistant

The AS212-C automotive socket features weather and vibration resistance in a durable nylon housing. It includes a right-angle bezel for placement under a panel, dashboard, or cabinet and features industry-standard 0.25×0.032-in. quick-fit male terminals that allow for easy, secure connections and fast terminations. Terminals also accept soldering directly to them with heavy-gauge wires. The

weather- and vibration-resistant cigarette-socket sets feature twist-and-lock mating for marine and automotive use. The AS212-C sells for $4.28 (1000).

Memory Protection Devices, www.memoryprotectiondevices.com

EDN ADVERTISER INDEX

Company	Page	Company	Page
Actel Corp	C-3	Lattice Semiconductor	4
Agilent Technologies	6, 35, 52	Linear Technology	19, 20, 44
austriamicrosystems AG	28	Maxim Integrated Products	47
Avnet Electronics Marketing	8, 25	Memory Protection Devices	36
Avtech Electrosystems Ltd	55	Mill-Max Manufacturing Corp	11
Digi-Key Corp	C-1, 3	Pico Electronics Inc	38, 39
EMA Design Automation	C-4	Renesas Technology Corp	49
Future Technology Devices	29	Samtec USA	15
Integrated Device Technology	C-2	Silicon Labs	23
Intel Corp	13	Stanford Research Systems Inc	27
International Rectifier Corp	7	Trilogy Design	55
Jameco Electronics	10	TTI Inc	17
Keil Software	37	Xilinx	31

EDN provides this index as an additional service. The publisher assumes no liability for errors or omissions.

Superman's X-ray vision would've come in handy

DANIEL VASCONCELLOS

Years ago, I was working as a manufacturing engineer on a team supporting production of digital medical X-ray machines. Each machine had an optional infrared remote control that paralleled the hardwired control panel for positioning the X-ray tube, the patient table, and other system components to make exposures. An outside vendor had designed the remote and the companion infrared receiver board.

Production testing of completed machines was causing the rejection of many infrared remote controls. Typically, the commanded function would initiate, but it would stop before completion. Strangely, some remote controls would perform better on one system than another or work well on one occasion but not another. I suspected timing issues on the remote control's embedded controller. The vendor insisted that it had tested each remote and that each had worked properly when it was shipped.

The infrared receiver board had redundant detectors. It also had a diagnostic port that showed how it was interpreting the coded commands from the remote controller. Pushing a button on the remote control generated a string of codes at the receiver board's diagnostic port. Examining this string, I saw that many of the characters were invalid, even when the system was correctly performing the requested function. The infrared board's proprietary software was obviously using an algorithm to decide which detector outputs and codes to use in generating a motion command.

The vendor's engineering staff was cooperative. The vendor confirmed that error-correcting algorithms in its proprietary software processed the received command characters and decided which to use. The company suggested some changes in the algorithm. We tried several changes, but they resulted in only slight improvements.

After much thought and experiment to analyze the cause of the problem, I finally asked the vendor for details of its testing process. Even though our system-design specifications were clear, I discovered that the vendor's routine testing of remote controls involved placing them on a fixture and verifying that pushing the buttons on the remote control generated the requested command outputs at a receiver board on the fixture.

This scenario differed greatly from our production environment, in which multiple operators tested completed systems, with only superficial control of the distance and aiming of the remote control. In addition, our manufacturing floor exposed the infrared link to wide variations in external illumination levels and multiple possible reflection paths, challenges typical of the environment of the system's ultimate deployment.

Once I explained to the vendor's software engineers that their tests were unrealistic, they began to tweak the system's processing of incoming commands. They quickly made software changes that improved the performance by more than an order of magnitude.

THE SIMPLEST EXPLANATION OF A PRODUCT PROBLEM IS SOMETIMES THE MOST LIKELY, BUT WATCH FOR EXCEPTIONS.

The lesson I learned was that the simplest explanation of a product problem is sometimes the most likely, but you should watch out for exceptions. In this case, it was essential to realize that this problem could involve more than just the remote controls. It was also crucial to question whether the vendor's test conditions reflected the product's end use, given that this end use—medical X-rays—clearly requires extreme precision. EDN

Reginald Neale is an engineer in Farmington, NY. His career spans a time line from vacuum tubes to embedded controllers.

www.edn.com/tales

SMARTFUSION™

"For us, it's more than a platform. It's a path to design freedom."

SmartFusion™— an FPGA, hard ARM® Cortex™-M3 and programmable analog on a single chip.

Integrating these technologies on a single platform gives you the freedom and flexibility to continually push your design further. Now you can optimize hardware/software tradeoffs at any point in development without incurring the cost and added time of board-level changes. With SmartFusion you can essentially create a microcontroller customized to the application rather than compromising with a processor that might fall short of your vision. *So go ahead and design fearlessly.*

www.actel.com/designfearlessly

Actel IS NOW Microsemi

OrCAD MAKES IT POSSIBLE
50% OFF UPGRADES*
EDIT AND VIEW FROM THE REVERSE SIDE

Wouldn't it be great if you could edit the bottom of a board without having to think backwards? OrCAD® makes it possible with the new Flip-Board feature in Cadence® OrCAD PCB Editor 16.3. Flip-Board is not just limited to viewing, as all design editing capabilities can be performed while in this mode. Don't let your tools limit the way you design. With just the click of a button you can easily view and edit designs on either side of the board without missing a beat.

Truly scalable PCB technology
Cadence OrCAD PCB design technologies comprise a complete, cost-effective PCB design solution, from design capture to final output. To successfully meet project goals, PCB designers and electrical engineers need powerful, intuitive, and integrated technologies that work seamlessly across the entire PCB design flow. OrCAD PCB design technologies offer fully integrated front-end design, analog/signal integrity simulation, and place-and-route technologies that boost productivity and shorten time to market.

***Save up to 50% today!**
This discounted pricing won't last long. All orders must be placed by December 17th to qualify. For more information contact EMA Design Automation, a Cadence Channel Partner, by phone: 800.813.7288, email: info@ema-eda.com, or online at www.ema-eda.com/discounts.

Scan with smartphone for movie

cadence®

EMA Design Automation™
Cadence Channel Partner